

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 233 624 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
21.08.2002 Bulletin 2002/34

(51) Int Cl.7: H04N 7/26

(21) Application number: 01129242.2

(22) Date of filing: 11.12.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

- Gormish, Michael J., c/o Ricoh Innovations, Inc.
Menlo Park, CA 94025-7022 (US)
- Boliek, Martin, c/o Ricoh Innovations, Inc.
Menlo Park, CA 94025-7022 (US)
- Wu, Gene K., c/o Ricoh Innovations, Inc.
Menlo Park, CA 94025-7022 (US)

(30) Priority: 15.02.2001 US 784928

(71) Applicant: Ricoh Company, Ltd.
Tokyo 143-8555 (JP)

(74) Representative: Schwabe - Sandmair - Marx
Stuntzstrasse 16
81677 München (DE)


(72) Inventors:
• Schwartz, Edward L., c/o Ricoh Innovations, Inc.
Menlo Park, CA 94025-7022 (US)

(54) A memory usage scheme for performing wavelet processing

(57) A system comprising a memory and wavelet processing logic is described. The memory is sized to include lines to store a band of an image and additional lines. The wavelet processing logic comprises a wavelet transform and access logic. The wavelet transform generates coefficients when applied to data in the memory. The access logic reads data from the memory into the

line buffers to supply data stored in the memory to the wavelet transform and to store coefficients in the memory, such that after data stored at a first pair of lines is read from memory into the buffers of the access logic. The access logic reuses the first pair of lines to store coefficients generated by the wavelet transform that are associated with a second pair of lines different from the first pair of lines.

FIG.2



EP 1 233 624 A1

Description**FIELD OF THE INVENTION**

- 5 [0001] The present invention relates to the field of compression and decompression; more particularly, the present invention relates to memory usage for performing wavelet processing.

BACKGROUND OF THE INVENTION

- 10 [0002] The new JPEG 2000 decoding standard (ITU-T Rec.T.800/ISO/IEC 15444.1:2000 JPEG 2000 Image Coding System) provides a new coding scheme and codestream definition for images. Although the JPEG 2000 standard is a decoding standard, the JPEG 2000 specifies encoding and decoding by defining what a decoder must do. Under the JPEG 2000 Standard, each image is divided into one or more rectangular tiles. If there is more than one tile, the tiling of the image creates tile-components that can be extracted or decoded independently of each other. Tile-components
15 comprise all of the samples of a given component in a tile. An image may have multiple components. Each of such components comprises a two-dimensional array of samples. For example, a color image might have red, green and blue components.

- [0003] After tiling of an image, the tile-components may be decomposed into different decomposition levels using a wavelet transformation. These decomposition levels contain a number of subbands populated with coefficients that
20 describe the horizontal and vertical spatial frequency characteristics of the original tile-components. The coefficients provide frequency information about a local area, rather than across the entire image. That is, a small number of coefficients completely describe a single sample. A decomposition level is related to the next decomposition level by a spatial factor of two, such that each successive decomposition level of the subbands has approximately half the horizontal resolution and half the vertical resolution of the previous decomposition level.

- 25 [0004] Although there are as many coefficients as there are samples, the information content tends to be concentrated in just a few coefficients. Through quantization, the information content of a large number of coefficients is further reduced. Additional processing by an entropy coder reduces the number of bits required to represent these quantized coefficients, sometimes significantly compared to the original image.

- [0005] The individual subbands of a tile-component are further divided into code-blocks. These code blocks can be
30 grouped into partitions. These rectangular arrays of coefficients can be extracted independently. The individual bit-planes of the coefficients in a code-block are entropy coded with three coding passes. Each of these coding passes collects contextual information about the bit-plane compressed image data.

- [0006] The bit stream compressed image data created from these coding passes is grouped in layers. Layers are
35 arbitrary groupings of successive coding passes from code-blocks. Although there is great flexibility in layering, the premise is that each successive layer contributes to a higher quality image. Subband coefficients at each resolution level are partitioned into rectangular areas called precincts.

- [0007] Packets are a fundamental unit of the compressed codestream. A packet contains compressed image data
from one layer of a precinct of one resolution level of one tile-component. These packets are placed in a defined order in the codestream.

- 40 [0008] The codestream relating to a tile, organized in packets, are arranged in one, or more, tile-parts. A tile-part header, comprised of a series of markers and marker segments, or tags, contains information about the various mechanisms and coding styles that are needed to locate, extract, decode, and reconstruct every tile-component. At the beginning of the entire codestream is a main header, comprised of markers and marker segments, that offers similar information as well as information about the original image.

- 45 [0009] The codestream is optionally wrapped in a file format that allows applications to interpret the meaning of, and other information about, the image. The file format may contain data besides the codestream.

- [0010] The decoding of a JPEG 2000 codestream is performed by reversing the order of the encoding steps. Figure
1 is a block diagram of the JPEG 2000 standard decoding scheme that operates on a compressed image data codestream. Referring to Figure 1, a bitstream initially is received by data ordering block 101 that regroups layers and
50 subband coefficients. Arithmetic coder 102 uses contextual information collected during encoding about the bit-plane compressed image data, and its internal state, to decode a compressed bit stream.

- [0011] After arithmetic decoding, the coefficients undergo bit modeling in coefficient bit modeling block 103. Next,
the codestream is quantized by quantization block 104, which may be quantizing based on a region of interest (ROI) as indicated by ROI block 105. After quantization, an inverse transform is applied to the remaining coefficients via
55 transform block 106, followed by DC and optional component transform block 107. This results in generation of a reconstructed image.

- [0012] The JPEG2000 standard leaves many choices to implementers.

SUMMARY OF THE INVENTION

[0013] A system comprising a memory and wavelet processing logic is described. The memory is sized to include lines to store a band of an image and additional lines. The wavelet processing logic comprises a wavelet transform and access logic. The wavelet transform generates coefficients when applied to data in the memory. The access logic reads data from the memory into the line buffers to supply data stored in the memory to the wavelet transform and to store coefficients in the memory, such that after data stored at a first pair of lines is read from memory into the buffers of the access logic. The access logic reuses the first pair of lines to store coefficients generated by the wavelet transform that are associated with a second pair of lines different from the first pair of lines.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0015] Figure 1 is a block diagram of the JPEG 2000 standard decoding scheme.

[0016] Figure 2 illustrates one embodiment of an organization for an image in memory.

[0017] Figures 3A-F illustrate a transform memory organization for various levels depicting conceptually how coefficients may be stored for the forward (Figures A-C) and inverse (Figures 3D-F) transforms.

[0018] Figures 4A and B illustrate embodiments of a single memory where the input image data and the various decomposition levels of the image data can be stored during forward and inverse transforms, respectively.

[0019] Figure 5 illustrates one embodiment of the process of handling the input data.

[0020] Figure 6A illustrates a system having a progression order conversion parser.

[0021] Figure 6B illustrates a progression converter converting from a resolution progressive codestream to a quality progressive codestream.

[0022] Figure 7A shows multiple ways to convert a codestream from one progression order to another.

[0023] Figure 7B shows one embodiment of simplified conversion paths to convert a codestream from one progression order to another.

[0024] Figure 8 illustrates one embodiment of a process for performing progression order conversion.

[0025] Figure 9 illustrates a decoder that selects portions of a codestream based on sideband information.

[0026] Figure 10 is a flow diagram of a process for using layers when decoding.

[0027] Figure 11 is a flow diagram of one embodiment of an editing process.

[0028] Figure 12 illustrates a bell-shaped curve of a range of values that are quantized to a particular value.

[0029] Figure 13 is a flow diagram of one embodiment of a process to reduce flicker.

[0030] Figure 14 illustrates one embodiment of an encoder (or portion thereof) that performs the quantization to reduce flicker.

[0031] Figure 15A illustrates a process for performing rate control.

[0032] Figure 15B illustrates an exemplary number of layers that may be subjected to first and second passes.

[0033] Figure 16 illustrates one embodiment of the process for accessing the groupings of tile parts.

[0034] Figure 17 and 18 illustrate quantizers for one component for a three level 5,3 transform.

[0035] Figure 19 illustrates an example of HVS weighted quantization.

[0036] Figure 20 is a block diagram of one embodiment of a computer system.

[0037] Figure 21 illustrates an example progression with tile parts for a single server.

[0038] Figure 22 illustrates an example of layering for a 5,3 irreversible transform.

[0039] Figure 23 illustrates an example in which transform has 5 levels and the data is divided up into layers 0-3.

[0040] Figure 24 illustrates one example of a situation in which flicker may be avoided in which values in first and third frames are used to set the value in the second frame.

[0041] Figure 25 is a block diagram of a prior art decoding process that includes color management.

[0042] Figure 26 illustrates one embodiment of a non-preferred camera encoder.

[0043] Figure 27 illustrates one embodiment of a simpler camera encoder.

[0044] Figure 28 is a flow diagram of one embodiment of a process for applying an inverse transform with clipping on partially transformed coefficients.

DETAILED DESCRIPTION OF THE PRESENT INVENTION

[0045] Improvements to compression and decompression schemes are described. It is a purpose of the techniques and implementations described herein to use choices in JPEG 2000 to make high speed, low cost, low memory and/or feature rich implementations.

[0046] In the following description, numerous details are set forth in order to provide a thorough explanation of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0047] Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0048] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0049] The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0050] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0051] A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"); random access memory ("RAM"); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

Overview

[0052] The following descriptions relate to implementations or novel ways to take advantage of the flexibility of JPEG 2000 or other coding schemes with similar features.

Memory Usage for Low Memory and Fast Burst Access

[0053] Figure 2 shows one embodiment of an organization for an image in memory 201. Referring to Figure 2, only the "tile height" raster lines, or a band of the image, are in memory 201, not the whole image. Thus, the amount of an image in memory 201 is equal to the image width multiplied by the tile height. Inside the band of the image is at least one tile, such as tile 210.

[0054] The wavelet transform processing logic 202 includes memory access logic 202A to read data from and store data to memory 201 to enable wavelet transform 202B to be applied to the data (image data or coefficients depending on the level of coefficient). Wavelet processing logic 202 may comprise hardware, software or a combination of both.

[0055] In one embodiment, access logic 202A accesses the tile with four parameters: a pointer or index to the start of the tile in memory, the width of the tile, the height of the tile, and the line offset to get from the start of one line to another (the image width). Alternatively, access logic 202A accesses memory 201 using a pointer or index to the end of the tile instead of the width of the tile.

[0056] In one embodiment, in order to access for each line of a tile or a portion of a line of an image to perform some function F, the following process may be used.

```

line = start
for y = 0 to tile_height - 1
  for x = 0 to tile_width - 1
    perform function F with line[x]
  line = line + line_offset

```

One of the functions F may include applying a wavelet transform on pairs of lines. Also another function F may be a DC level shift, multiple component transform.

[0057] Such a process would be performed by processing logic that may comprise hardware (e.g., dedicated logic, circuitry, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both.

[0058] In one embodiment, coefficients from a subband are accessed using a similar scheme with a starting point, width, height and line offset. Because rows of coefficients are stored together in memory, rows may be accessed efficiently when stored in cache, burst accessible memories or memories that are wider than one coefficient.

[0059] Figures 3A-C show a transform memory organization for various levels depicting conceptually how coefficients may be stored. All LH, HL and HH coefficients (using the nomenclature of ITU-T Rec.T.800/ISO/IEC 15444:2000 JPEG 2000 Image Coding System) are coded. These coefficients are represented by dotted lines in Figures 3B and C. Input lines of input tile 301 and LL coefficients (shown as solid lines in Figures 3B and 3C in successive levels) only need to be stored temporarily while computing the transform with the exception of the final transform level's LL coefficients which are coded. When a transform is used that does the horizontal and vertical transforms in one pass and uses line buffers, once a pair of input rows has been completely read (input lines or LL coefficients), the space used by lines can be reused.

[0060] Figures 3A-C show input tile 301, level 1 (L1)(302) and level 2 (L2)(303) memory areas aligned with an offset to indicate how reuse might be accomplished in one embodiment. The addition of two rows, rows 312 and 313, to the memory space used to hold input tile 301, is needed to generate the L1 coefficients when reusing the memory for input tile 301 for L1 coefficients. The addition of four rows, rows 341-342, to the memory space used to hold the L1 coefficients is needed to generate the L2 coefficients when reusing the memory storing the L1 coefficients for L2 coefficients. (Note that there are two rows between rows 341 and 342 that are wasted space.) The additional lines are preferably behind the direction the wavelet transform is being applied to the information in the memory.

[0061] More specifically, a pair of input rows input tile 301 corresponds to one row of each of LL, LH, HL and HH coefficients at level 1, resulting from the application of a transform to two different lines and the results of applying the wavelet transform being written into lines of the memory. For example, the results of applying a wavelet transform to input rows 310 and 311 are the coefficients in portions of rows 312 and 313 of L1 coefficients (302). For example, LL coefficients 321 of row 312 corresponds to the LL coefficients (solid line) of level 1, HL coefficients 322 of row 312 corresponds to the HL coefficients of level 1, LH portion 323 of row 313 corresponds to the LH coefficients of level 1, and HH portion 324 corresponds to the HH coefficients of level 1. Note that the level 1 coefficients from the first two input lines are stored in two extra rows at the top of the memory with the remaining level 1 coefficients being written into the locations storing the data of input tile 301 to reuse the memory. The width and height for each type of coefficient (e.g., LH, HL, HH) for each subband of level 1 coefficients is half that of input tile 301. The line offset from the LL row to the next LL row for level 1 (e.g., the distance from row 312 to row 314 in Figure 3B) is twice the tile width (since each coefficient row is from an area corresponding to two lines).

[0062] Similarly, the results of applying the wavelet transform to two rows of LL coefficients at level 1 (solid lines) are the coefficients in two rows namely LL coefficients (331), LH coefficients (332), HL coefficients (333) and HH coefficients (334) at level 2. The width and height for level 2 coefficients is a quarter that of input tile 301. The line offset for level 2 is four times the tile width (since each coefficient row is from an area corresponding to two level 1 LL rows or four input lines). Thus, four extra lines of memory are needed to use the same memory that is storing the input tile to store the L2 coefficients. Note that if a third decomposition level was being performed, an additional 8 lines would be needed. Thus, in this example, a total of 14 extra lines are needed to enable reuse of the memory that stores an input tile and has two levels of decomposition applied thereto. A general formula may be used to determine the number of extra lines is as follows:

$$2^{(\text{maxlevel}+1)} \cdot 2.$$

[0063] To access subbands, such as the LL, LH, HL and HH subbands, only a starting pointer and the offset between

rows/lines are necessary. The height and width are also needed to know when to stop when accessing a tile.

[0064] As the number of decomposition levels increases, some rows at the bottom of memory become unused. That is, the lines of memory below the L1 coefficients after the first decomposition level become unused, the lines of memory below the L2 coefficients after the second decomposition level become unused, etc. In one embodiment, this extra space may be reused.

[0065] Figures 3D-3F illustrate the corresponding inverse transform memory usage in which additional lines store the results of applying an inverse transform and those additional lines are in the memory behind the direction the inverse transform is being performed.

[0066] Figure 4A shows one embodiment of a single memory where the input and the various levels can be stored during application of a forward transform. Referring to Figure 4A, locations for the input tile, level 1 coefficients, level 2 coefficients, and level 3 coefficients is shown with the added 2, 4 and 8 lines respectively. Figure 4B shows a similar single memory embodiment where the input coefficients of various levels of the transform can be stored along with the output during application of an inverse transform.

[0067] Table 1 shows the amount of memory required for various transform levels for a 256x256 tile for separate memories and reused memory.

Table 1

level	Separate memory (bytes)		reused memory (bytes)	
1	256x256 =	65,536	2x256 =	512
2	128x128 =	16,384	4x256 =	1,024
3	64x64 =	4,096	8x256 =	2,048
4	32x32 =	1,024	16x256 =	4,096
5	16x16 =	256	32x256 =	8,192
6	8x8 =	64	64x256 =	16,384

[0068] For reused memory, the amount listed is the additional new memory used for that level. For this example, reusing memory for levels 1, 2 and 3 saves memory. Level 4 may use a separate memory.

[0069] The memory for levels 4, 5 and 6 could be placed in a single memory after level 3 has been generated or in a completely different and separate memory. The amount of memory necessary is 38x32, which is less than 5x256. Because there are two unused lines after generating the level 1 coefficients (i.e., the memory that stored the last two lines of input data), a small memory savings can be achieved by letting the levels 4, 5 and 6 reuse these two lines. This is particularly important because the number of additional lines for levels 4, 5, and 6 is 16, 32 and 64, and the extra space between the lines will be twice as far and half as wide as the level before.

[0070] In one embodiment, coefficients from levels 4, 5, and 6 are packed in a smaller memory structure, such as storage area 450 in Figure 4. Referring to Figure 4, the level 4 coefficients are stored in an area having a height equal to the tile height divided by 8 (2^3 where 3 corresponds to the number of levels) and a width equal to the tile width w divided by 8 (2^3 where 3 corresponds to the number of levels previously stored elsewhere). An additional two lines 451 are all that is needed to store level 5 coefficients in the same necessary storage area. Similarly, an additional four lines is all that is necessary to accommodate using this memory storage area for the level 6 coefficients. Note that no lines are skipped when storing the coefficients. In one embodiment in which a 256x256 tile is being processed, the extra 5 lines at the bottom of storage area 430, two lines 421 and approximately 4.75 lines 422 are used to accommodate storage area 450. As shown, the approximate by three lines 422 represent allocated memory or in addition to that necessary to store the input tile. In this manner, the storage area for the input tile is almost completely reused.

[0071] In one embodiment, to use a very little, or potentially minimum, memory, level 6 is stored separately from levels 4 and 5. However, this only saves 64 bytes of memory.

[0072] A memory a little smaller than 273x256 can hold all the transform coefficients for a 256x256 tile. This is less than 7% more than a true in-place memory organization. Unlike an in-place memory organization, extra copies are avoided while simultaneously keeping the rows packed together for fast access.

[0073] Table 2 shows another example of using separate versus reused memory for 128x128 tiles. For this size, the first three transform levels can reuse memory in a 142x128 buffer.

Table 2

level	Separate memory (bytes)		reused memory (bytes)	
1	128x128 =	16,384	2x128 =	256
2	64x64 =	4,096	4x128 =	512

Table 2 (continued)

level	Separate memory (bytes)		reused memory (bytes)	
3	32x32=	1,024	8x128=	1024

[0074] In one embodiment, a decision to use in-place memory or new memory is a function of tile height and transform level. Such a decision may be based on the following:

if tile height $> 2^{(3 \cdot \text{level} - 2)}$, then use in-place method

if tile height $= 2^{(3 \cdot \text{level} - 2)}$, then either may be used

if tile height $< 2^{(3 \cdot \text{level} - 2)}$, then use new memory

To illustrate the application of the decision, Table 3 below:

Table 3

level	$2^{(3 \cdot \text{level} - 2)}$
1	2
2	16
3	128
4	1024
5	8192

[0075] In some applications, adapting the memory organization to the tile height is inconvenient. A single fixed memory organization can be used. Tile sizes smaller than 128x128 typically result in bad compression performance, so would typically not be used. While tile sizes bigger than 1Kx1K can be used for very large images, this does not significantly improve compression and the large amount of memory required would typically be burdensome. Therefore, assuming a tile height between 128 and 1024 inclusive and using in-place memory for 3 levels of the transform is a good heuristic.

[0076] Decoding is similar in that the results of applying an inverse transform are written ahead of where the decoding processing logic is reading, with the only notable difference being that the start is from the highest level to the lowest level, such as level 6 to level 1 in the example above. In such a case, the input tile ends up at the top of the memory structure. The extra lines to accommodate the memory reuse are in decreasing order. For example, using the structure of Figure 4B, 8 lines would be necessary to create the L2 coefficients from the L3 coefficients, 4 extra lines would be necessary to create the L1 coefficients from the L2 coefficients and 2 extra lines would be necessary to create the input tile from the L1 coefficients.

[0077] In one embodiment, to handle input tile data, a color conversion may be performed on the data prior to encoding. Figure 5 illustrates one embodiment of the process of handling the input data. Referring to Figure 5, color input pixels are received in raster order. These color pixels may be in RGB, YCrCb, CMY, CMYK, grayscale, etc. The color input pixels may be stored as tiles in a memory, such as memory 501, by band (or other forms).

[0078] Pixels from storage 501 or received directly from the input undergo color conversion and/or level shifting, with the resulting outputs being stored in one coefficient buffers 502₁-502_N. That is, once the color conversion has been completed on each tile, it is stored in one of the coefficient buffers 502₁-502_N, and then the next tile can be processed. In one embodiment, there is one coefficient buffer for each component.

[0079] Coefficient buffers 502₁-502_N are used by the transform in the manner described above to perform the wavelet transform while reusing memory. Thus, coefficient buffers 502₁-502_N are both input and output to wavelet transform.

[0080] After the transform is applied to coefficient buffers 502₁-502_N, the context model 503 and entropy coder 505 can perform further compression processing on the already transformed data. The coded data is buffered in coded data memory 505.

[0081] While performing the further compression processing on one tile, the transform may be applied to another tile. Similarly, any or all the operations may be performed on multiple tiles at the same time.

Progression Order Conversion

[0082] In the JPEG2000 standard, data in a compressed codestream can be stored in one of the five progression orders. The progression order can change at different points in the codestream. The order is defined by embedded "for layers" on layers, precincts, resolution, and components.

[0083] Five progression orders are described in the standard in Table A-16 of the JPEG 2000 standard. They are layer-resolution-component-position progression (LRCP), resolution-layer-component-position progression (RLCP), resolution-position-component-layer progression (RPCL), position-component-resolution-layer progression (PCRL), component-position-resolution-layer progression (CPRL).

[0084] The order may be defined in the COD or POC markers of the JPEG 2000 standard. The Coding style default (COD) marker is defined by the JPEG 2000 standard and describes the coding style, number of decomposition levels, and layering that is the default used for compressing all components of an image (if in the main header) or a tile (if in a tile-part header). The Progression order change (POC) marker describes the bounds and progression order for any progression order other than that specified in the COD marker segments in the codestream. The Packet Length Main Header (PLM) indicates a list of packet lengths in tile-parts for every tile part in order and the Packet Length, Tile-part header (PLT) indicates tile packet lengths in a tile-part and indicates where the data is in the codestream.

[0085] The JPEG 2000 standard in section B.12 only specifies how packets of compress data are formed for a given progression order. It does not describe how data should be converted from one progression order to another progression order.

[0086] In one embodiment, a progression order converting parser converts a codestream to a desired progression order based on the user input without decoding the data and then encoding it again. Figure 6A illustrates a system having such a parser. Referring to Figure 6A, parser 601 receives requests from a client for a particular progression order. The client may be viewing a web page and selects a particular link. In response to the request, parser 601 accesses server 602 to obtain the codestream associated with full image 603 from memory 604 and converts the codestream into a different progression order based on the request. The request indicates the progression order by using an optional command (e.g., RL2L (Resolution-layer progression to Layer Progression)). The progression order that is described may be based on layer, resolution, component, precinct, or tile.

[0087] Figure 6B illustrates the progression converter converting from a layer progressive codestream (LRCP) to a resolution progressive (RLCP) codestream. The progression orders map directly to each other.

[0088] Figure 7A shows multiple ways to convert a codestream from one progression order to another. Referring to Figure 7A, each of the five progressions (LRCP, RLCP, RPCL, CPRL, and PCRL) are shown with paths to each of the others, such that all progressions are shown. In one embodiment, the parser causes all conversions to go through the layer progression first and then to a selected conversion. Figure 7B shows one embodiment of such simplified conversion paths in which the number of required mappings is reduced from 10 (as in Figure 7A) to 4. However, any one of the five progression orders could be used as the one to which all are converted before arriving at the selected order. The conversion technique described herein simplifies source codes in that the number of lines of source code is much less than the multiple ways of conversion. This results in less debug time and fewer memory and run-time variables.

[0089] To perform the conversion, the order of the packets in the codestream must be reordered. The packets are labeled by their sequential order in the codestream. Markers may indicate the starting point of the data, the length of the data (or alternatively the endpoint of the data) and how the data should be handled. For example, the indication of how the data is to be handled may indicate whether the data is to be deleted, whether the data is to be truncated, or some other operation to be performed on the data. Such handling information may also come from rate distortion information, such as may be provided in a PLT/PLM and/or the PPT/PPM marker sets of the JPEG 2000 standard. In this manner, the codestream may be truncated without changing the packet header.

[0090] In one embodiment, a list, array, or other structure (such as reordering structure 601A) is built by indicating the portion of data in each packet. Using this structure, the packets may be reordered.

[0091] Figure 8 illustrates one embodiment of a process for performing progression order conversion. The process is performed by processing logic that may comprise hardware (e.g., dedicated logic, circuitry, etc.), software (such as is run by, for example, a general purpose computer or dedicated machine), or a combination of both.

[0092] Referring to Figure 8, the process begins by processing logic building a list from headers in the packets (processing block 801) and optionally marking list items "delete" for quantization (processing block 802). Next, processing logic reorders the list to map the original progression to a desired progression (including handling input and output with progressions specified with POC markers (bounds on the progression order) (processing block 803). Thereafter, processing logic outputs coded data based on reordered list (processing block 804).

[0093] Therefore, the combination of re-ordering and parsing allows specification of the desired ordering and resolution, quality, etc.

A Progression Order Conversion Example

[0094] The following is an example showing how packets are arranged in a codestream. The codestream was formed based on 2 components, 2 layers, 3 decomposition levels, and layer progression.

[0095] Table 4 shows the packet order, length and association index of packets in the example. The packet order column shows the sequential order of packets placed in a codestream. The length indicates the length of the packets. The association index shows the resolution, layer, component, and precinct of the packet.

[0096] For example, packet[0] is the first packet in the codestream after the first tile header. It has a length of 589 bytes. Association index R0L0C0P0 indicates the packet belongs to resolution w, layer x, component y and precinct z.

Table 4

Packet order	Length	Association Index
packet[0]	length=589	R0L0C0P0
packet[1]	length=589	R0L0C1P0
packet[2]	length=924	R1L0C0P0
packet[3]	length=924	R1L0C1P0
packet[4]	length=1602	R2L0C0P0
packet[5]	length=1602	R2L0C1P0
packet[6]	length=733	R3L0C0P0
packet[7]	length=733	R3L0C0P0
packet[8]	length=535	R0L1C0P0
packet[9]	length=535	R0L1C1P0
packet[10]	length=1523	R1L1C0P0
packet[11]	length=1523	R1L1C1P0
packet[12]	length=5422	R2L1C0P0
packet[13]	length=5422	R2L1C1P0
packet[14]	length=16468	R3L1C0P0
packet[15]	length=16468	R3L1C1P0

[0097] In this codestream, packets are grouped based on the layer in which they reside. The first 8 packets belong to Layer 0. The following 8 packets belong to Layer 1.

[0098] Using the conversion process described herein, the above codestream is converted to resolution layer progression. The following shows how the above packets are re-ordered.

[0099] After the layer progressive codestream is converted to resolution progression, in the new codestream, packets are grouped based on resolution. Such a grouping is shown in Table 5. The first 4 packets belong to resolution 0, the next 4 packets to resolution 1, and so on.

Table 5

Previous Packet order	Packet order	Length	Association Index
0	packet[0]	length=589	R0L0C0P0
1	packet[1]	length=589	R0L0C1P0
8	packet[2]	length=535	R0L1C0P0
9	packet[3]	length=535	R0L1C1P0
2	packet[4]	length=924	R1L0C0P0
3	packet[5]	length=924	R1L0C1P0
10	packet[6]	length=1523	R1L1C0P0
11	packet[7]	length=1523	R1L1C1P0
4	packet[8]	length=1602	R2L0C0P0
5	packet[9]	length=1602	R2L0C1P0
12	packet[10]	length=5422	R2L1C0P0

Table 5 (continued)

Previous Packet order	Packet order	Length	Association Index
13	packet[11]	length=5422	R2L1C1P0
6	packet[12]	length=733	R3L0C0P0
7	packet[13]	length=733	R3L0C1P0
14	packet[14]	length=16468	R3L1C0P0
15	packet[15]	length=16468	R3L1C1P0

*One Embodiment of a Conversion Algorithm***[0100]** Resolution to Layer Progression

```

n = 0;
for(l=0;l<layer;l++){
  for(r=0;r<resolution+1;r++){
    for(c=0;c<component;c++){
      new_packet[n] = old_packet[l*component + r*layer*component +
c];      n++;
    }
  }
}

```

Layer to Resolution Progression

```

n = 0;
for(r=0;r<resolution+1;r++){
  for(l=0;l<layer;l++){
    for(c=0;c<component;c++){
      new_packet[n] = old_packet[r*component +
l*(resolution+1)*component + c];
      n++;
    }
  }
}

```

where layer = the number of layers in a codestream,
resolution = the number of decomposition levels in a codestream, and
component = the number of components in a codestream

Data Hiding (Sideband Information) in JPEG2000 Coding

[0101] Bit hiding allows sideband information to be transmitted without increasing the file size. Sideband information that does increase file size but does not break naive decoders might also be valuable (although the COM marker defined by the JPEG 2000 standard might be used instead).

[0102] Some marker segments, packet headers and packets are padded out to the nearest byte. Examples of the JPEG 2000 marker segments include PPM, PPT, PLM, and PLT. In addition, some marker segments can be longer than they need to be including QCD, QCC, and POC. In all of these cases, the padded data values are not defined.

[0103] Several proprietary coding schemes could use this semi-randomly located undefined data to provide a number of important types of information including, but not limited to, decoding and filtering hints, ownership, segmentation hints, and so on. A hint might include an index to a particular enhancement scheme. For example, if it is known that an image is mostly text, a value may be sent that indicates that a first post-processing filter is to be used. On the other hand, if the area is mostly a graphic image, then a value may be sent that indicates that a second post-processing filter is to be used.

[0104] The following are places where bits may be hidden or sideband information may be stored in the codestream.

- arithmetic coder (AC) termination (without predictable termination)
- end of packet header rounding to byte
- after last packet, before next tile
- tag tree construction by not always using minimum
- packet header Lblock signalling
- LSB parity for codeblocks (refinement pass only, cleanup pass only, all)
- QCD, QCC extra subbands, POC.

[0105] For example, with respect to hiding data using AC termination, 0 to 7 bits are provided, at least, everytime the coder is terminated. However, this could be extended for a few bytes. These extra bits and bytes may be used for sending extra information.

[0106] With respect to each packet header, the end of a packet header is rounded to a byte boundary. Therefore, there may be 1 to 7 bits that may be available for sending extra information at times when rounding would have been necessary. Similarly, each packet is rounded to a byte boundary, thereby providing 1 to 7 bits (assuming that rounding would have been necessary). Also the last packet in a tile-part can be extended a few bytes. These extra bytes may be used to send additional information.

[0107] The length of the compressed data for a code-block can be given in the packet header with a non-minimum representation. The choice of representation (e.g., a non-minimum representation) could be used for indicating other information.

[0108] With respect to tag tree data hiding, packet headers of the JPEG 2000 standard use tag trees for coding first inclusion and zero bitplane information. When there are multiple codeblocks, tag trees are like a quadtree of minimum values. For example, in the case of 16 codeblocks in a 4x4 arrangement in a packet, the arrangement may be as follows:

10	7	12	15
3	20	21	5
81	45	5	9
18	8	12	24

An example tag tree, which is minimal for the 4x4 arrangement above is as follows:

3	0	2	7	4	7	10
	5	2	0	17	16	0
			73	37	0	4
			10	0	7	19

in which "3" is added to every codeblock's value, and "0", "2", "5" and "2" are each added to the 4 corresponding codeblocks. Finally, there is one value per codeblock. That is, the minimal tag tree is created by taking the first 2x2 group in the 4x4 arrangement above and look at minimum value is out of the four values. In this case, for the 2x2 block

10	7
3	20

the minimum value is 3. This is then performed on the other 2x2 blocks. Then these identified minimum values are evaluated again to determine their minimum, which would be "3" in the example. Then the minimum value is subtracted from the four minimum values to create the following

0	2
5	2

Then, for the remaining numbers in the 4x4, the number 3 is subtracted from each value along with the value in the 2x2 that corresponds to the particular value in the 4x4 arrangement, thereby resulting in the tag tree above.
The first row adds up as follows:

$$10=3+0+7$$

$$7=3+0+4$$

$$12=3+2+7$$

$$15=3+2+10$$

[0109] A variable length code may be used that efficiently represents small numbers.

[0110] An example of a tag tree that is not minimal is as follows:

2	1 3	7	4	7	10
	6 3	0	1 /	16	0
		73	37	0	4
		10	0	7	19
(Note that representing "3", "0", "2", "5" and "2" might use less bitstream data than "2", "1", "3", "6" and "3".)					

[0111] Once a tag tree representation has been made, a determination can be made as to whether the representation is minimal or not based on whether there is a zero in the 2x2 block. Therefore, this information is hidden. For example, the 1 bit block represents the 1 in the 2x2 block above indicates it is not part of a minimal tag tree, but can be used to convey some particular information to a decoder. Likewise if a 2 was the minimal value in the 2x2 block, such a fact may convey different information to a decoder.

[0112] The JPEG 2000 POC, QCD, and QCC markers can have redundant entries. It is as if the codestream were quantized and the markers were not rewritten. For example, the QCD and QCC markers have values for a number of subbands specified by the syntax of the marker. If there are fewer subbands actually coded in the bitstream, data may be hidden in the values used for the missing subbands. The redundant entries may be replaced and used for hidden or sideband information.

[0113] The hidden or sideband information may include post-processing hints (such as, for example, sharpen this tile with a specified filter or strength, or smooth, or perform optical character recognition (OCR) on this region, etc.), decoding hints, security (such as, for example, an encryption key for decoding the remainder of the image or another image, etc.) codestream identification (such as, for example, labeling POTUS as the originator of the file, etc.) and/or other information.

Use of Layers When Encoding

[0114] Layers are part of the JPEG standard. In one embodiment, sideband information, possibly in a COM marker, is used by the decoder to allow selecting of layers during decoding. The sideband information may be used to select layers for postcompression quantization to meet rate/distortion targets for different viewing distances, different resolutions, different regions of interest, different frequency content for analysis (e.g., finding edges of text).

[0115] In one embodiment, the layers are predefined based on rate. For example, the first layer represents a 1-bit per pixel image, while the second layer represents a 2-bit per pixel image, etc. Therefore, the layers run from the lowest quality to the highest quality. Likewise, target rates can be met for lower resolutions as well.

[0116] The sideband information may be stored in a marker segment of the codestream. In one embodiment, the JPEG 2000 comment (COM) marker is used to provide information about the layers. Specifically, the COM marker may

EP 1 233 624 A1

be used to indicate the number of bytes for each resolution and/or rate across the entire image or a relative number of bytes for each additional layer. Table 6 indicates each layer and its resolution in the number of bytes across the tile in an image. Such a table may have distortion values instead.

Table 6

5	lev=0	layer=0	comp=0	bytes=529
	lev=0	layer=0	comp=1	bytes=555
	lev=0	layer=0	comp=2	bytes=493
10	lev=0	layer=1	comp=0	bytes=129
	lev=0	layer=1	comp=1	bytes=130
	lev=0	layer=1	comp=2	bytes=123
	lev=0	layer=2	comp=0	bytes=7
	lev=0	layer=2	comp=1	bytes=8
15	lev=0	layer=2	comp=2	bytes=12
	lev=0	layer=3	comp=0	bytes=1
	lev=0	layer=3	comp=1	bytes=1
	lev=0	layer=3	comp=2	bytes=129
20	lev=1	layer=0	comp=0	bytes=705
	lev=1	layer=0	comp=1	bytes=898
	lev=1	layer=0	comp=2	bytes=712
	lev=1	layer=1	comp=0	bytes=146
	lev=1	layer=1	comp=1	bytes=114
25	lev=1	layer=1	comp=2	bytes=116
	lev=1	layer=2	comp=0	bytes=224
	lev=1	layer=2	comp=1	bytes=250
	lev=1	layer=2	comp=2	bytes=263
30	lev=1	layer=3	comp=0	bytes=201
	lev=1	layer=3	comp=1	bytes=212
	lev=1	layer=3	comp=2	bytes=200
35	lev=2	layer=0	comp=0	bytes=889
	lev=2	layer=0	comp=1	bytes=1332
	lev=2	layer=0	comp=2	bytes=1048
	lev=2	layer=1	comp=0	bytes=240
	lev=2	layer=1	comp=1	bytes=329
	lev=2	layer=1	comp=2	bytes=328
40	lev=2	layer=2	comp=0	bytes=599
	lev=2	layer=2	comp=1	bytes=767
	lev=2	layer=2	comp=2	bytes=725
	lev=2	layer=3	comp=0	bytes=335
	lev=2	layer=3	comp=1	bytes=396
45	lev=2	layer=3	comp=2	bytes=420
50	lev=3	layer=0	comp=0	bytes=1
	lev=3	layer=0	comp=1	bytes=395
	lev=3	layer=0	comp=2	bytes=402
	lev=3	layer=1	comp=0	bytes=251
	lev=3	layer=1	comp=1	bytes=450
	lev=3	layer=1	comp=2	bytes=562
	lev=3	layer=2	comp=0	bytes=525
	lev=3	layer=2	comp=1	bytes=990
55	lev=3	layer=2	comp=2	bytes=1313
	lev=3	layer=3	comp=0	bytes=1214

EP 1 233 624 A1

Table 6 (continued)

lev=3	layer=3	comp=1	bytes=1798
lev=3	layer=3	comp=2	bytes=2585

[0117] In another embodiment, the ordering could be by layer. Thus, the information above is consolidated for each level (not segregated by level or component), as shown below:

Ordering by layer=0 bytes=7959 bitrate=0.971558 PSNR=30.7785

Ordering by layer=1 bytes=10877 bitrate=1.327759 PSNR=32.0779

Ordering by layer=2 by tes= 16560 bitrate=2.021484 PSNR=35.7321

[0118] Distortion by layers can be based on PSNR. For example,

layer=0 PSNR=30.7785

layer=1 PSNR=32.0779

layer=2 PSNR=35.7321

[0119] In an alternative embodiment, such information may be hidden in the codestream as described above. The information may be used to control rate distortion.

[0120] In another embodiment, the layers may be predefined for a particular viewing distance. In such a case, the data is divided into layers from the highest frequency, lowest resolution to the lowest frequency, highest resolution.

[0121] In one embodiment, the layer information indicates the summation of bits across the entire image for that layer and all previous layers (for example the 16,011 bits listed next to layer 1 indicates the total number of bits for layer 0 and layer 1). Alternatively, bytes, words, kilobytes, or other units of memory or rate could be used instead of bits. Table 7 shows this type of absolute rate information.

[0122] Table 8 shows relative rate information. Layer 0 has 4096 bits, layer 1 has 11,915 bits, etc.

Table 7

layer	Rate (bytes)
0	4,096
1	16,011
2	40,000
3	100,000
4	250,000
5	500,000
6	1,000,000
7	2,500,000
8	5,500,000

Table 8

layer	Rate (bytes)
0	4,096
1	11,915
2	23,989
3	60,000
4	150,000
5	250,000
6	500,000
7	1,500,000
8	3,000,000

[0123] For example, if only 750,000 bytes may be allowed in the decoded image, then all that can be decoded (as

the 1,000,000 bytes tabulated with layer 6 includes the 500,000 bytes of layers 0-5) is through layer 5 and half of importance layer 6. In some embodiments, no packets from layer 6 would be included. In other embodiments, some packets from layer 6 would be included and others would be replaced by zero packets so that the total amount of layer 6 data was approximately 250,000 bytes.

[0124] Figure 22 illustrates an example of layering for a 5,3 irreversible transform with three levels, MSE or similar. Referring to Figure 22, there are 45 layers shown. Each additional layer improves MSE in an order that gives good rate-distortion for MSE.

[0125] Figure 23 illustrates another example in which transform has 5 levels and the data is divided up into layers 0-3. Layer 0 corresponds to the thumbnail version, layers 0-1 correspond to the monitor (or screen) resolution, layers 0-2 correspond to the print resolution, and layers 0-3 correspond to lossless.

[0126] In an alternative embodiment, the layers may be predefined for some other distortion metric (e.g., MSE, weighted MSE, sharpness of text, etc.)

[0127] The decoder uses the information regarding the layers from the codestream to select layers to generate an image. The decoder knowing what the desired viewing characteristics from the application or implementation (see Table 9 below), and using the information from the codestream specifying the layers, can quantize the codestream in order to display an image at the correct viewing distance. Figure 9 illustrates such a decoder. Referring to Figure 9, decoder 901 receives a codestream and includes quantization logic 902 that examines the COM marker and uses information about the viewing distance it is at stored in storage 903 to generate quantized codestream 904 via, for example, selecting the proper layers. Quantized codestream 904 is decoded by decoding logic 905 (e.g., a JPEG 2000 decoder) after selecting layers to generate an image data 906. A naive decoder would simply ignore the data in the comment marker.

[0128] Figure 10 is a flow diagram of a process for using layers when decoding. The process is performed by processing logic that may comprise hardware (e.g., dedicated logic, circuitry, etc.), software (such as is run by, for example, a general purpose computer or a dedicated machine), or a combination of both.

[0129] Referring to Figure 10, the process begins by processing logic receiving a codestream of compressed logic data (processing block 1001). The image data is organized into multiple layers, each of which comprises coded data that adds visual value to the image (e.g., look sharper, better defined, better contrast, etc.). Next processing logic selects one or more layers for quantization based on sideband information (processing block 1002). After selection, processing logic decompresses the non-quantized layers of the codestream (processing block 1003).

Editing of Tiles, Tile-parts, or Packets

[0130] Once a codestream is created, it may be desirable to edit parts of the image. That is, for example, after performing encoding to create the codestream, a set of tiles may be decoded. After decoding the set of tiles, editing may be performed, followed by encoding the set of tiles with the edits to the same size as the encoded tiles were prior to their decoding. Examples of typical editing include sharpening of text and removing "red-eye." The JPEG 2000 codestream can be edited in memory or in a disk file system without rewriting the entire codestream.

[0131] Figure 11 is a flow diagram of one embodiment of an editing process. The process is performed by processing logic that may comprise hardware (e.g., dedicated logic, circuitry, etc.), software (such as is run by, for example, a general purpose computer or a dedicated machine), or a combination of both.

[0132] Referring to Figure 11, processing logic initially determines the tiles, tile-parts, or packets that cover the area, resolution, components, and/or precincts to be edited and decodes them (processing block 1101). This determination may be made in response to a user selecting an area and/or working resolution. The determination may use editing information for a higher resolution to determine which parts or tiles cover the portion to be edited. Once decoding has been completed, processing logic performs the desired edits (processing block 1102).

[0133] After performing the desired edits, processing logic recompresses the data into coded data (processing block 1103) and creates a replacement tile, tile-part, or packet for the codestream (processing block 1104). In one embodiment, in creating the replacement tile, tile-part, or packet, processing logic pads out the data with bytes at the end of the codestream if the new data is smaller than the unedited version of the data to make the replacement tile, tile-part or packet the same size as the unedited version.

[0134] In an alternative embodiment, processing logic may use a marker, or tag, such as a COM marker segment of the appropriate length instead of the padding. The COM marker could be used to fill space or could contain information that the encoder wanted to include. It could contain information such as, for example, sideband information described herein or a copyright license for an image or text or other file format information.

[0135] In one embodiment, in creating the replacement tile, tile-part, or packet, processing logic truncates the last packets for any or all components until the data fits in the codestream if the new data is larger than the unedited version of the data.

[0136] Editing of an image may be performed by changing coded data for tiles, tile-parts, or codeblocks. In one

Optimal Encoder Quantization

[0139] In some cases, two spatially adjacent coefficients may be close to each other numerically yet in separate quantization bins, such as coefficient values 1201 of range R_2 and 1202 of range R_1 in Figure 12. The results of the quantization may cause an artifact to occur. In one embodiment, for coefficients near a boundary between two quantization bins, the encoder selects a bin such as Range R_1 into which a coefficient, such as coefficient 1201, will be quantized so that it is consistent with neighbors, such as coefficient 1202. This helps avoid artifacts. That is, this technique reduces distortion yet may increase rate, particularly when a coefficient is moved from a smaller bin to a higher bin.

[0140] At times, flicker occurs when applying wavelet compression to motion sequences. An example of such flicker may include the image getting brighter or darker in areas or the appearance of edges changing in successive frames as the motion sequence is played (mosquito noise around the edges). The flicker may be due to the application of different local quantization to successive frames of a motion sequence or to noise exacerbated by quantization that is viewed temporarily.

[0142] A test of whether to apply such quantization to a coefficient value in a subsequent frame is based on the quantization that was performed on the coefficient in the previous frame. Thus, the encoder is utilizing frame dependency to eliminate flicker while the decoder decodes data frame by frame independently.

if $|Q(D_{n+1}) - (D'_n)| < T$

$$D'n+1 = D'n$$

else

$$D'n+1 = Q(Dn+1)$$

For example, the value T may be twice the quantization step size. Other values of T include, but are not limited to, $\sqrt{2}Q$, $1.5Q$, $2\sqrt{2}Q$.

[0144] One of the coefficient values may be modified to be either a predetermined closeness to another coefficient value. The closeness may be determined by some threshold. The threshold may be user set or adaptive based on some criteria. The threshold could be different based on the subband and, perhaps, on the persistence of the particular value (number of frames that this coefficient is close). In one embodiment, the coefficient value is set equal to the other coefficient value. In alternative embodiments, the coefficient is set to be within the quantization bin size of the other coefficient value or twice the quantization bin size.

[0145] Figure 14 illustrates one embodiment of an encoder (or portion thereof) that performs the quantization described above. Referring to Figure 14, a quantizer 1400 receives coefficients 1410 for frames of a motion sequence from a wavelet transform (not shown). The coefficients are received by quantization logic 1401 which compares a threshold value stored in memory 1401 to coefficient values for the previous frame that are stored in memory 1403 to coefficients 1410 with a scalar quantizer Q applied from memory 1404.

[0146] Quantization logic 1401 may comprise comparison hardware (e.g., logic with gates, circuitry, etc.) or software to perform the comparison. This comparison hardware and software may implement a subtractor or subtraction operation. The results are a quantized codestream (assuming some values have been changed.)

[0147] This may be applied over two or more frames. Also the comparison is not limited to two consecutive frames. The comparison can be over 3, 4, 5, etc., frames, for example, to determine if a variance exists. Figure 24 illustrates one example in which values in a first and third frame are used to set the value in the second frame.

[0148] Note that the quantization can also be codestream quantization with a code block-based rule.

Rate Control, Quantization, and Layering

[0149] In one embodiment, selective quantization of coefficients can be performed during encoding by setting a subset of the refinement bits to be the more probable symbol (MPS). This may be performed at a user selected bitplane. For examples, if there is text on a background image, with a goal of having sharp text images while minimizing coded data required for the background, the refinement bits that are set to MPS are those that do not effect text for the last bitplane, while using the actual value for bits that effect text.

[0150] Such a quantization scheme may be used to implement non-uniform quantization step sizes. For example, if one wanted to have a background with fewer bits, setting the refinement bits to the MPS could operate as a form of quantization. This quantization scheme causes some level of distortion but lowers the bit rate necessary to transfer the codestream.

[0151] Note that although this technique may be applied to bits generated during the refinement pass, the technique has application to other compression schemes (e.g., lists generated during subordinate passes, tail bits of CREW of Ricoh Silicon Valley, Menlo Park, California, MPEG IV texture mode, etc.).

[0152] In one embodiment, the same technique may be applied to other changes between frames. That is, in one embodiment, a change due to a rate distortion in one frame may be performed in a subsequent frame to avoid distortion effects.

Rate Control and Quantization

[0153] In one embodiment, user specified quantization is provided. For a 3 level transform for one component, 7 quantization values are sufficient: level 1 HH, level 1 HL and LH, level 2 HH, level 2 HL and LH, level 3 HH, level 3 HL and LH, and level 3 LH.

[0154] If quantization values are bitplanes to truncate (which is equivalent to scalar quantization by powers of 2), 3-bit values (0...7) are sufficient for most applications. (For image components with depth 12-bits or more and 5 or more transform levels, perhaps higher quantizations might be useful.) Values 0...6 could be used to specify the number of bitplanes to truncate and 7 could be used to mean discard all bitplanes. The three bit values may be written to a controller that controls compression (or decompression) hardware (e.g., JPEG2000 compatible hardware) to perform the quantization.

[0155] For 3 component color quantization:

- 21 values can be used with separate values for each component,
- 14 values can be used, 7 for luminance and 7 for chrominance,
- 5 • 17 values can be used for 4:1:1 subsampled data, 7 for luminance and 5 for each chrominance component,
- 12 values can be used for 4:1:1 subsampled data, 7 for luminance and 5 for chrominance,
- 19 values can be used for 4:2:2 subsampled data, 7 for luminance and 6 for each chrominance component, and
- 13 values can be used for 4:2:2 subsampled data, 7 for luminance and 6 for chrominance.

10 Since $21 \times 3 = 63$ bits is less than 8 bytes, transferring or storing the quantization uses little resources. A central processing unit (CPU) might select one predetermined quantizer from a table and write it to a CPU or other controller controlling special purpose JPEG 2000 hardware (a chip) for each frame of a motion JPEG 2000 video sequence. Alternatively, one implementation of JPEG 2000 might have a small memory that holds 8 or 16 different quantizers that could be selected for each frame.

15 [0156] Quantizers can also be used to assign bitplanes to layers. For example, Q_0 , Q_1 , and Q_2 may be quantizers that specify bitplanes of coding pass to quantize. Quantizer Q_0 causes the most loss, while quantizer Q_2 causes the least loss. Layer 1 is all the data quantized by Q_0 but not quantized by Q_1 . Layer 2 is all the data quantized by Q_1 but not quantized by Q_2 . Layer 3 is all the data quantized by Q_2 .

20 *Simple Quantization*

[0157] Figures 17 and 18 show example quantizers (label A...Q) for the 3-level 5/3 transform as the number of coefficient LSBs to truncate or not code. Truncating N bitplanes is equivalent to a scalar quantizer of 2^N . The subband where the quantization changes with respect to the previous quantizer is highlighted with a dashed box. The quantizers D, K and Q all have the same relationship between the subbands. Other quantizers might be used that are better for MSE or for other distortion metrics.

25 [0158] The exemplary Verilog below converts a single quantization value "q" into seven quantizers (number of LSBs to truncate). The variable q_1_HH is used for level 1 HH coefficients, the variable q_1_H is used for level 1 HL and LH coefficients, etc. Some consecutive values of q result in the same quantizer: 0 and 1; 2 and 3; 4 and 5; $8i+6$ and $8i+7$ for all integers i with $i \geq 0$.

35

40

45

50

55

```

module makeQ(q, q_1HH, q_1H, q_2HH, q_2H, q_3HH, q_3H,
q_3LL);

```

```

5      input [5:0] q;
      output [3:0] q_1HH;
      output [3:0] q_1H;
      output [3:0] q_2HH;
10     output [2:0] q_2H;
      output [2:0] q_3HH;
      output [2:0] q_3H;
      output [2:0] q_3LL;

15     wire [3:0] temp_2H;
      wire [3:0] temp_3HH;
      wire [3:0] temp_3H;
      wire [3:0] temp_3LL;
20     wire [2:0] qlo;
      wire [2:0] qhi;

      assign qlo = q[2:0];
      assign qhi = q[5:3];

      assign q_1HH    = qhi + ((qlo >= 2) ? 1 : 0);
      assign q_1H     = qhi + ((qlo >= 4) ? 1 : 0);
30     assign q_2HH    = qhi + ((qlo >= 6) ? 1 : 0);
      assign temp_2H   = qhi + ((qlo >= 1) ? 0 : -1);
      assign temp_3HH  = qhi + ((qlo >= 3) ? 0 : -1);
      assign temp_3H   = qhi + ((qlo >= 5) ? 0 : -1);
35     assign temp_3LL  = qhi - 1

      assign q_2H     = (temp_2H < 0) ? 0 : temp_2H;
      assign q_3HH    = (temp_3HH < 0) ? 0 : temp_3HH;
40     assign q_3H     = (temp_3H < 0) ? 0 : temp_3H;
      assign q_3LL    = (temp_3LL < 0) ? 0 : temp_3LL;

```

```

endmodule

```

Human Visual System Weighting for Color and Frequency

[0159] Table 9 shows additional bitplanes to quantize (e.g., truncate) for luminance to take advantage of the frequency response of the Human Visual System (from Table J-2 of the JPEG 2000 standard). A viewing distance of 1000 pixels might be appropriate for viewing images on a computer monitor. Larger viewing distances might be appropriate for print images or television.

Table 9 -

Human Visual System Weighting for Luminance			
subband	extra biplanes to quantize for viewing distance of ...		
	1000 pixels	2000 pixels	4000 pixels
1HH	2	4 or 5	discard all
1HL, 1LH	1	2 or 3	6
2HH	—	2	4 or 5
2HL, 2LH	—	1	2 or 3
3HH	—	—	2
3HL, 3LH	—	—	1

Additionally chrominance may be quantized more heavily than luminance.

[0160] Figure 19 shows a quantization that starts with Figure 17(D) and then adds frequency weighting for a 1000 pixel viewing distance (to both luminance and chrominance), keeps 3LL chrominance unchanged, discards 1HL and 1HH chrominance for 4:2:2 and additional 2 bitplanes are discarded for the remaining chrominance.

[0161] Sharp text without ringing artifacts is more desirable than exact gray value for text/background. That is, if a gray level is supposed to be at 50% (for example), and is instead at 60%, it is often not visually objectionable if the image is of text. In one embodiment, the LL (DC) coefficients are quantized more heavily for text than for non-text images at low bitrate. For example, for an 8-bit image component, a quantization step size of 8, 16 or 32 might be used for text only regions and a quantization step size of 1, 2 or 4 might be used for regions containing non-text. This allows more fidelity for the high frequency coefficients, thereby resulting in text with sharp edges.

Using Quantizers to Divide Things into Layers

[0162] Table 10 shows 16 example quantizers. Quantizer 15 is lossless. Quantizer 8 is the same as Figure 19. These can be used divide the subband bitplanes into layers.

Table 10

subband	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Y 1HH	all	all	6	6	5	5	4	4	3	3	2	2	1	1	0	0
Y 1HL, LH	6	5	5	4	4	3	3	2	2	1	1	0	0	0	0	0
Y 2HH	5	4	4	3	3	2	2	1	1	0	0	0	0	0	0	0
Y 2HL, LH	4	4	3	3	2	2	1	1	0	0	0	0	0	0	0	0
Y 3HH	4	4	3	3	2	2	1	1	0	0	0	0	0	0	0	0
Y 3HL, LH	4	3	3	2	2	1	1	0	0	0	0	0	0	0	0	0
Y 3LL	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
C ₁ 1HL, HH	HL and HH always discarded for 4:1:1 or 4:2:2 only															
C ₁ 1LH	all	all	all	all	6	6	5	5	4	4	3	3	2	2	1	0
C ₁ 2HH	all	6	6	5	5	4	4	3	3	2	2	1	1	0	0	0
C ₁ 2HL, LH	all	6	5	5	4	4	3	3	2	2	1	1	0	0	0	0
C ₁ 3HH	all	6	5	5	4	4	3	3	2	2	1	1	0	0	0	0
C ₁ 3HL, LH	all	5	5	4	4	3	3	2	2	1	1	0	0	0	0	0
C ₁ 3LL	all	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
C ₂ 1HL, HH	HL and HH always discarded for 4:1:1 or 4:2:2 only															

Table 10 (continued)

subband	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
C ₂ 1LH	all	all	all	all	6	6	5	5	4	4	3	3	2	2	1	0
C ₂ 2HH	all	6	6	5	5	4	4	3	3	2	2	1	1	0	0	0
C ₂ 2HL, LH	all	6	5	5	4	4	3	3	2	2	1	1	0	0	0	0
C ₂ 3HH	all	6	5	5	4	4	3	3	2	2	1	1	0	0	0	0
C ₂ 3HL, LH	all	5	5	4	4	3	3	2	2	1	1	0	0	0	0	0
C, 3LL	all	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

[0163] Layer 0 contains all data not quantized away by quantizer 0. This would be luminance data only: all of 3LL; all but 4 bitplanes of 2HL, 2LH, 3HL, 3LH and 3HH; all but 5 bitplanes of 2HH and all but 6 bitplanes of 1HL and 1LH. Layer 1 contains all data not in layer 0 and not quantized away by quantizer 1. This would be luminance bitplanes 5 for 1HL and 1LH, bitplane 4 for 2 HH, bitplane 3 for 3HL and 3LH; all 3LL chrominance; all but 5 bitplanes for chrominance 3HL and 3LH; and all but 6 bitplanes for chrominance 2HL, 2LH and 3HH. Finally, layer 15 would contain the LSB of 1LH chrominance.

Rate Control with Multiple Layers and Tile-Parts

[0164] There several well known techniques for rate control in compression systems. The simplest method is to pick a distortion for every "unit" compressed (a unit may be an 8x8 block in JPEG, a frame in a motion sequence, a tile of a single image, a subband of a tile in a wavelet coded image, etc.). If the distortion selected leads to a bitrate higher than the desired average bitrate, the distortion allowed is increased for new units as they are compressed. If the distortion selected leads to a bit rate lower than the desired average bitrate, the distortion allowed is decreased for new units as they are compressed.

[0165] A more complex method buffers the compressed data from some number of "units." The bitrate and/or distortion for each unit at each distortion level is stored. Then the distortion to allow across all the units in the buffer is determined when the buffer is full. If the buffer is sufficient to contain the entire image, extremely high quality results can be obtained. In JPEG 2000, layers are designed to contain increments to quality. Thus, selecting a distortion can mean selecting the number of layers to use for each code block or tile. A complete description of this type of rate control is in, David Taubman, "High Performance Scalable Image Compression with EBCOT," IEEE Transactions on Image Processing.

[0166] There are several disadvantages to this process. One disadvantage is that a buffer memory for the entire codestream is required. A second disadvantage is that the latency (time until any of the codestream is output) is high. A third disadvantage is that the second pass could take large amount of time.

[0167] To mitigate these problems, each tile of a JPEG 2000 codestream is encoded as described above with at least two layers. At the completion of encoding each tile, a number of packets (e.g., layer, resolution, precinct, tile-component) are output to the codestream as a complete tile-part. The remaining layers are stored in the buffer. A second pass through the remaining coded data in the buffer is optional. During this second pass, extra packets from each tile are appended to the codestream as complete tile-parts as space or time allows. If in a fixed-rate application, then only packets within the given rate are appended. If in a fixed time application, then only number of cycles allowed. One embodiment of this process is shown in Figure 15A. Thus, these can be the 2 complete tile-parts output for each tile.

[0168] Figure 15B illustrates a number of layers, layers 1-n. Layer 1 is output on the first pass, and the remaining layers are most likely below fixed-time or fixed-rate time limits. Layer 2 may be output on a second pass within fixed-time or fixed-rate requirements while achieving similar distortion over all the components.

[0169] The above process is advantageous in that it allows the buffer to store a fraction of the coded data, the first data can be output (transmitted or stored) sooner, and the second pass through the data can be faster because there is less data to process. Also less memory is required for buffering.

[0170] The criterion for selecting which packets go into the first set of tile-parts can be similar to any other rate control algorithm. In one embodiment, the rate of packets can be less than the desired average bitrate for the whole image. For example, if a final compressed bitstream at 2.0 bpp is desired, the first pass could place 1.5 bpp for every tile in the codestream, and buffer 1 bpp for every tile.

[0171] The second pass can select from the remaining data the packets to place in the second tile part of each tile. Thus, to obtain a 2.0 bpp average encoding, some tiles that had high distortion after the first pass could receive all the remaining data saved for the tile, while other tile parts which had low distortion after the first pass might not have any

additional data transmitted.

Rate Control for Compressed Codestream Data

5 **[0172]** Some rate control techniques described herein include rate control performed on a compressed codestream based on a request implemented by selecting some number of layers to keep in the codestream. A parser may be used to produce a new codestream which shows the bitrate based on layers. This bitrate is equal to or less than the bitrate specified by the request.

10 **[0173]** The parser may use a data structure referred to herein as a "packet structure." Note that this data structure may be used for other purposes such as, for example, the versatile pocket data structure described below. In one embodiment, the packet structure includes a packet start pointer and packet length. It also contains a tile number, a resolution, a component, layer, and a precinct the packet belongs to. Finally, it also consists of a selection flag. This flag, when set to a predetermined value (e.g., 1), indicates if the packet is selected in the array for writing out to a new codestream.

15 **[0174]** In one embodiment, packets are read in sequential order from a codestream based on the progression order information indicated by the COD marker.

20 **[0175]** The number of bytes is computed based on the bitrate desired by the request. The number of bytes belonging to layer 0 is added up to a total. Then this total of bytes is compared with the number of bytes desired. If the total is less than the number of bytes desired, one additional layer is added to the total. The process continues until the total is equal to or greater than the number of bytes desired or all packets have been added.

25 **[0176]** During the process, those packets which have been added to the total, are marked as selected by the selection flag in the structure.

30 **[0177]** If the total is equal to the number of bytes desired, the addition process is stopped. If the total exceeds the number of bytes desired, the packets in the last layer added are subtracted from the total. This is done to guarantee that the bitrate is below the bitrate desired. Consequently, during the subtraction step, packets which have been subtracted from the total are marked unselected.

35 **[0178]** In one embodiment, the related markers such as SOT, COD, PLT are updated according to the request. Packets are written to the new codestream. The packet structure may be created using the following:

30

35

40

45

50

55

```

typedef struct _PACK_ {          /* packet structure */
    int      start; /* packet starting point */
    int      length; /* packet length */
    unsigned short t; /* tile number the packet belongs to */
    unsigned short r; /* resolution the packet belongs to */
    unsigned short c; /* component the packet belongs to */
    unsigned short l; /* layer the packet belongs to */
    unsigned short p; /* precinct the packet belongs to */
    unsigned char  select; /* selection flag */
} Pack_t;

/* Store packets from tp->tile[i].Size[j] array to the packet structure array
**/* Layer progression (LRCP) order */

if(progression_order == 0){
    j = 0;
    for(i=0;i<number_of_tile;i++){
        m = 0;
        for(l=0;l<layer;l++){
            for(r=0;r<resolution+1;r++){
                for(c=0;c<component;c++){
                    for(p=0;p<precinct[r];p++){

                        tp->pk[j].start = tp->tile[i].pointer[m];
                        tp->pk[j].length = tp->tile[i].Size[m];
                        total_length += tp->tile[i].Size[m];

                        tp->pk[j].t = i;
                        tp->pk[j].r = r;

                        tp->pk[j].l = l;
                        tp->pk[j].c = c;
                        tp->pk[j].p = p;
                        m++;
                        j++;
                    }
                }
            }
        }
        num_packet[i] = m;
    }
}

```

Versatile Packet Data Structure

[0179] The same packet data structure described above can be used to facilitate other parsing options, once packets are read into the structure.

[0180] For resolution parsing, the packets which are to be excluded are marked unselected. For example, given a 4 resolution codestream, and a request is to produce a 3-resolution codestream, a parser marks all packets which belong to resolution 4 unselected. Then the newly produced codestream contains only packets from resolution 1 up to resolution 3.

[0181] Similarly, for component parsing, progression conversion parsing, quality parsing can be performed step by step processing the packets in the structure.

[0182] The packet data structure can handle complex requests. For example, a request which requires the parser to produce a codestream which has a 3-resolution, 2-layer, and 1-component codestream.

Clipping After Each Inverse Transform

[0183] As a result of quantization performed on wavelet coefficients, the final decoded pixels are often outside of the original range of allowed pixels from the specified bit depth. Typically, these pixels are clipped to the original range so that further image processing or display devices can use the original bit depth.

[0184] For example, an eight bit image has pixel values between 0 and 255, inclusive. After lossy compression is used, the decoded image may contain values like -5 and 256. To provide an eight bit output, these values are clipped to 0 and 255 respectively. This clipping procedure always reduces pixel wise distortion because the original image did not contain pixels outside of the clipping bounds. This procedure is well known and recommended by the JPEG 2000 standard.

[0185] In addition to the bounds on the final output samples, there are bounds on the values coefficients can assume at the various stages of the wavelet transform. Just as quantization can change the final decoded samples to lie outside the original bounds, quantization can change the partially transformed wavelet coefficients to lie outside their original bounds. If these coefficients are clipped to their original bounds, distortion will decrease.

[0186] For example, after a horizontal (one dimensional) 5-3 reversible transform as specified by JPEG 2000 with 8 bit input samples, the maximum value of the low pass coefficient is +191, and the minimum possible value is -191. The high pass coefficient must be between -255 and 255 inclusive. After the vertical one dimensional transform, the Low-Low coefficients are bounded by -286 and 287. Thus when decoding an eight bit image, when the first level low-low pass coefficients are generated (by the inverse wavelet transform from a higher level), the coefficients can be clipped to -286 and +287, and distortion will decrease. Likewise after the first level vertical inverse transformation is done, the low pass coefficients can be clipped to -191, +191, and the high pass coefficients can be clipped to -255, 255.

[0187] For each subband, each filter, each transform level, and each image depth, there is a different maximum and minimum value for the coefficients. These maximum and minimum values can be computed by finding the signal that leads to the maximum and minimum and running the forward compression system and recording the maxima. The signals that lead to extreme values come from inputs where each pixel is either a maximum or minimum. Which pixels should be maximum and which pixels should be minimum can be determined by convolving sequences which are -1 when the wavelet coefficient is negative and +1 when the wavelet coefficient is positive. For the 5-3 filter used in JPEG 2000 Part I, the low pass signal of interest is [-1+1+1+1-1] and the high pass signal is [-1+1-1-1+1].

[0188] The signal (image) which will generate the largest LL value is:

+1	-1	-1	-1	+1
-1	+1	+1	+1	-1
-1	+1	+1	+1	-1
-1	+1	+1	+1	-1
+1	-1	-1	-1	+1

(where +1 must be replaced by the input maximum (e.g., 255) and -1 must be replaced by the input minimum (e.g., 0).

[0189] For irreversible filters, it is not necessary to actually run the system to determine the maxima, simply convolving the wavelet coefficients is sufficient. For the reversible 5-3 filter, however, the floor function is used in the computation of coefficients and is also used to determine the correct maxima.

[0190] Note that this may be used for other filters (e.g., a 9-7 filter).

[0191] Figure 28 is a flow diagram of one embodiment of a process for applying an inverse transform with clipping on partially transformed coefficients. The process is performed by processing logic, which may comprise hardware (e.g., circuitry, dedicated logic, etc.), software (such as that which runs on a general purpose computer system or a

dedicated machine), or a combination of both.

[0192] Referring to Figure 28, processing logic applies a first level inverse transform to coefficients (processing block 2801). Thereafter, processing logic clips the partially transformed coefficients to a predetermined range (processing block 2802). Next, processing logic applies a first level inverse transform to the clipped coefficients (processing block 2803) and clips the partially transformed coefficients to a predetermined range (processing block 2804), which is different than the range in processing block 2802. Again, processing logic applies a first level inverse transform to clipped coefficients (processing block 2805) and clips the partially transformed coefficients to still another predetermined range (processing block 2806).

10 Simplified Colorspace Handling

[0193] A typical decoding process including color management is shown in Figure 25. Referring to Figure 25, a file with a file format (e.g., a file format described in the JPEG 2000 standard) containing a restricted ICC profile is provided to a decoding device. Decompression block 2501 decompresses the file by taking the codestream portion of the file and performing context modeling, entropy decoding, and applying an inverse wavelet transform, but does not perform color space operations. If the codestream indicates the RCT or ICT component transform should be used to decode the codestream, these will be performed by block 2502. That is, inverse RCT/ICT block 2502 takes the components and the "RCT Y/N" indication (RCT if yes, ICT is no) and performs the specified inverse transform and provides (non-display) RGB pixels. (If specified by the syntax, inverse level shifting is also performed.)

[0194] Finally, the ICC color profile from the file format along with information about the display device will be used to produce the output pixels.

[0195] Inverse ICC block 2503 receives the (non-display) RGB pixels and the ICC profile and applies an inverse color space transform to provide display RGB pixels.

[0196] Figure 26 illustrates one embodiment of a non-preferred camera encoder. Referring to Figure 26, a camera generates YCrCb pixels. A converter 2602 converts the YCrCb pixels to RGB pixels and provides those two a typical JPEG 2000 encoder. The encoder comprises a RCT to ICT converter 2603 followed by a compressor 2604. The compressor generates an ICC_A for codestream.

[0197] Figure 27 illustrates one embodiment of a simpler camera encoder. That is, instead of including RCT/ICT converter 2603 and compressor 2604, a simple camera encoder includes only compressor block 2702. Referring to Figure 27, a camera 2701 generates YCrCb pixels and provides them to compressor 2702. Compressor comprises a JPEG 2000 encoder without an RCT conversion and generates an ICC_B codestream with RCT equaling 1 (with syntax signaling that the inverse RCT should be used on decoding). The relationship between ICC_B and ICC_A is given by the following equation:

$$ICC_B = ICC_A \circ YCrCb^{-1} \circ RCT$$

where \circ represents function composition.

[0198] Restricted ICC profiles are "syntaxes" for functions on pixels. A camera will typically write the same profile for all images, so ICC_B is computed offline, and copied into each output file. In a prior art system there must be HW for YCrCb⁻¹ and RCT/ICT which operates on every pixel.

Coding 4:2:2 and 4:1:1 Data as 4:4:4 Data with Quantization

[0199] The JPEG 2000 standard is typically used to handling data in a 4:4:4 format. It is not capable of describing how to reconstruct data in 4:1:1 or 4:2:2 formats in a 4:4:4 format for output. In one embodiment, when encoding 4:1:1 data, the encoder treats 1 HL, 1 LH and 1 HH coefficients as zero. When encoding 4:2:2 data, the encoder treats 1 HL and 1 HH coefficients as zero. Thus, with all information in the extra subbands quantized to zero, a decoder is able to receive the codestream in a way it expects. In other words, the encoded data resembles 4:4:4 data that has been heavily quantized.

File Order for Thumbnail, Monitor, Printer, and Full Resolution and Quality

[0200] Multiple images at multiple resolutions are important in many image processing situations. Depending on the application, a user may want to select different images of different resolutions. For example, thumbnail images may be used as an index into a large number of images. Also, a screen resolution image may be the image used to send to a monitor for display thereon. A print resolution image may be of lower quality for printer applications.

[0201] In one embodiment, a codestream of an image is organized into sections so that different versions of the

image, such as, for example, a thumbnail version, a screen version, a print version and a lossless version, is progressive by quality.

[0202] In one embodiment, the packets are arranged such that certain packets correspond to particular resolutions such as a thumbnail. The combination of these packets with other packets represents the monitor resolution image, which when combined with other packets may represent the printer version, etc. Using the POC and tile parts, portions of a codestream may be grouped together. For example, all the tiles of the thumbnail size may be grouped together followed by tiles for another resolution followed by tiles of another resolution, etc. Figure 21 illustrates an example progression with tile parts for a single server. Each tile's thumbnail is grouped in tile-parts at the beginning of a file. Figure 21A illustrates that tile-part 2101 is the only portion that is used for a thumbnail image. Figure 21B illustrates that for a monitor resolution, tile-parts 2102-2104 have been included with tile-part 2101. Figure 21C illustrates that for a printer resolution, tile-parts 2105 and 2106 have been included with tile-parts 2101-2104. Lastly, Figure 21D illustrates that for a lossless version of the data, the remaining three tile-parts 2107-2108 are included with the rest of the tile-parts. These sets of tile parts may be placed on a server in this progressive order.

[0203] One embodiment of the process for accessing the groupings of tile parts is shown in Figure 16. The process may be performed by processing logic that may comprise hardware (e.g., dedicated logic, circuitry, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. The following steps assume that the image has been transformed with sufficient resolution levels and layers to divide the image into the four sizes.

[0204] Referring to Figure 16, processing logic initially determines the correct resolution and layering for the thumbnail (processing block 1601). In one embodiment, to determine the correct resolution and layering for the thumbnail, processing logic creates a POC constrained to that resolution and layer for each tile and then creates a set of tile-parts and places this POC for each tile in the codestream.

[0205] Next, processing logic repeats processing block 1601 for the monitor resolution given that the thumbnail packets are already in the codestream (processing block 1602). Then, processing logic repeats processing block 1601 for the printer resolution given that the monitor packets are already in the codestream (processing block 1603).

[0206] Lastly, processing logic creates a POC marker with the extremes of the resolutions and layers for each tile (processing block 1604). In one embodiment, creating the POC with the extremes of the resolutions and layers is performed by creating a fourth set of tile-parts with the remaining tile-parts for a lossless version.

[0207] Note that the particular orders of the packets defined in the POCs are not of importance, only the limits.

An Exemplary Computer System

[0208] Figure 20 is a block diagram of an exemplary computer system. Referring to Figure 20, computer system 2000 may comprise an exemplary client 150 or server 100 computer system. Computer system 2000 comprises a communication mechanism or bus 2011 for communicating information, and a processor 2012 coupled with bus 2011 for processing information. Processor 2012 includes a microprocessor, but is not limited to a microprocessor, such as, for example, Pentium™, PowerPC™, Alpha™, etc.

[0209] System 2000 further comprises a random access memory (RAM), or other dynamic storage device 2004 (referred to as main memory) coupled to bus 2011 for storing information and instructions to be executed by processor 2012. Main memory 2004 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor 2012.

[0210] Computer system 2000 also comprises a read only memory (ROM) and/or other static storage device 2006 coupled to bus 2011 for storing static information and instructions for processor 2012, and a data storage device 2007, such as a magnetic disk or optical disk and its corresponding disk drive. Data storage device 2007 is coupled to bus 2011 for storing information and instructions.

[0211] Computer system 2000 may further be coupled to a display device 2021, such as a cathode ray tube (CRT) or liquid crystal display (LCD), coupled to bus 2011 for displaying information to a computer user. An alphanumeric input device 2022, including alphanumeric and other keys, may also be coupled to bus 2011 for communicating information and command selections to processor 2012. An additional user input device is cursor control 2023, such as a mouse, trackball, trackpad, stylus, or cursor direction keys, coupled to bus 2011 for communicating direction information and command selections to processor 2012, and for controlling cursor movement on display 2021.

[0212] Another device that may be coupled to bus 2011 is hard copy device 2024, which may be used for printing instructions, data, or other information on a medium such as paper, film, or similar types of media. Furthermore, a sound recording and playback device, such as a speaker and/or microphone may optionally be coupled to bus 2011 for audio interfacing with computer system 2000. Another device that may be coupled to bus 2011 is a wired/wireless communication capability 2025 to communication to a phone or handheld palm device.

[0213] Note that any or all of the components of system 2000 and associated hardware may be used in the present invention. However, it can be appreciated that other configurations of the computer system may include some or all of

the devices.

[0214] Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as essential to the invention.

Claims

1. A system comprising:

a memory sized to include lines to store a band of an image and additional lines;
a wavelet processing logic comprising

a wavelet transform to generate coefficients when applied to data in the memory;
access logic to read data from the memory into the line buffers to supply data stored in the memory to the wavelet transform and to store coefficients in the memory, such that after data stored at a first pair of lines is read from memory into the buffers of the access logic, the access logic reuses the first pair of lines to store coefficients generated by the wavelet transform that are associated with a second pair of lines different from the first pair of lines.

2. The system defined in Claim 1 wherein the access logic stores coefficients in contiguous lines of memory with coefficients from the same subband and decomposition level adjacent each other.

3. The system defined in Claim 1 wherein a first line of each of the first and second pairs of lines are located in the memory at an offset with respect to each other.

4. The system defined in Claim 3 wherein the access logic stores the first outputs of the wavelet transform for each coefficient level in the additional lines within a distance of the offset.

5. The system defined in Claim 3 wherein size of the offset is different for each transform level.

6. The system defined in Claim 3 wherein the size of the offset is equal to:
 $2^{(transform\ level\ of\ coefficient\ being\ stored)}$

7. The system defined in Claim 6 wherein, during decomposition, the offset for storing the first rows of each pair of rows of L1 coefficients in the memory is two lines from the first row of data of the image associated with said each pair of rows of the L1 coefficients, and the offset for storing the first row of each pair of rows of L2 coefficients is four lines from the first row of L1 coefficients associated with said each pair of rows of the L2 coefficients.

8. The system defined in Claim 1 wherein the access logic stores coefficients associated with a decomposition level greater than level three in the lines of the memory that previously stored the band of the image.

9. The system defined in Claim 3 wherein the addition lines relating to the offset are above the line storing the band of the image.

10. The system defined in Claim 1 wherein the wavelet transform is a forward wavelet transform.

11. The system defined in Claim 1 wherein the wavelet transform is an inverse wavelet transform.

12. A method comprising:

reading data from a memory into line buffers to apply a wavelet transform thereto; and
storing coefficients created by applying the wavelet transform at lines in the memory so that each set of coefficients generated from data stored at each pair of lines in the memory is stored in the memory at lines that are at an offset with respect to said each pair of lines in the memory.

13. The method defined in Claim 12 further comprising access logic reusing a first pair of lines to store coefficients generated by a wavelet transform, that are associated with a second pair of lines different from the first pair of lines, after data stored at a first pair of lines is read from memory into the buffers of the access logic, and wherein a first line of each of the first and second pairs of lines are located in the memory at an offset with respect to each other.
14. The method defined in Claim 13 further comprising the access logic storing the first outputs of the wavelet transform for each coefficient level in additional lines within a distance of the offset.
15. The method defined in Claim 13 wherein size of the offset is different for each transform level.
16. The method defined in Claim 13 wherein the size of the offset is equal to:
 $2^{(\text{transform level of coefficient being stored})}$.
17. The method defined in Claim 16 wherein, during decomposition, the offset for storing the first rows of each pair of rows of L1 coefficients in the memory is two lines from the first row of data of the image associated with said each pair of rows of the L1 coefficients, and the offset for storing the first row of each pair of rows of L2 coefficients is four lines from the first row of L1 coefficients associated with said each pair of rows of the L2 coefficients.
18. The method defined in Claim 12 further comprising access logic storing coefficients associated with a decomposition level greater than level three in the lines of the memory that previously stored the band of the image.
19. The method defined in Claim 13 wherein the addition lines relating to the offset are above the line storing the band of the image.
20. An article of manufacture comprising at least one recordable media storing executable instructions thereon which, when executed by a processing device, cause the processing device to:

read data from a memory into line buffers to apply a wavelet transform thereto; and
store coefficients created by applying the wavelet transform at lines in the memory so that each set of coefficients generated from data stored at each pair of lines in the memory is stored in the memory at lines that are at an offset with request to said each pair of lines in the memory.
21. The article of manufacture defined in Claim 20 further comprising instructions, which when executed by the processing device cause the processing device to reuse a first pair of lines to store coefficients generated by a wavelet transform, that are associated with a second pair of lines different from the first pair of lines, after data stored at a first pair of lines is read from memory into the buffers of the access logic, and wherein a first line of each of the first and second pairs of lines are located in the memory at an offset with respect to each other.
22. An apparatus comprising:

means for reading data from a memory into line buffers to apply a wavelet transform thereto; and
means for storing coefficients created by applying the wavelet transform at lines in the memory so that each set of coefficients generated from data stored at each pair of lines in the memory is stored in the memory at lines that are at an offset with request to said each pair of lines in the memory.

FIG.1

Reconstructed and rendered image




FIG.2

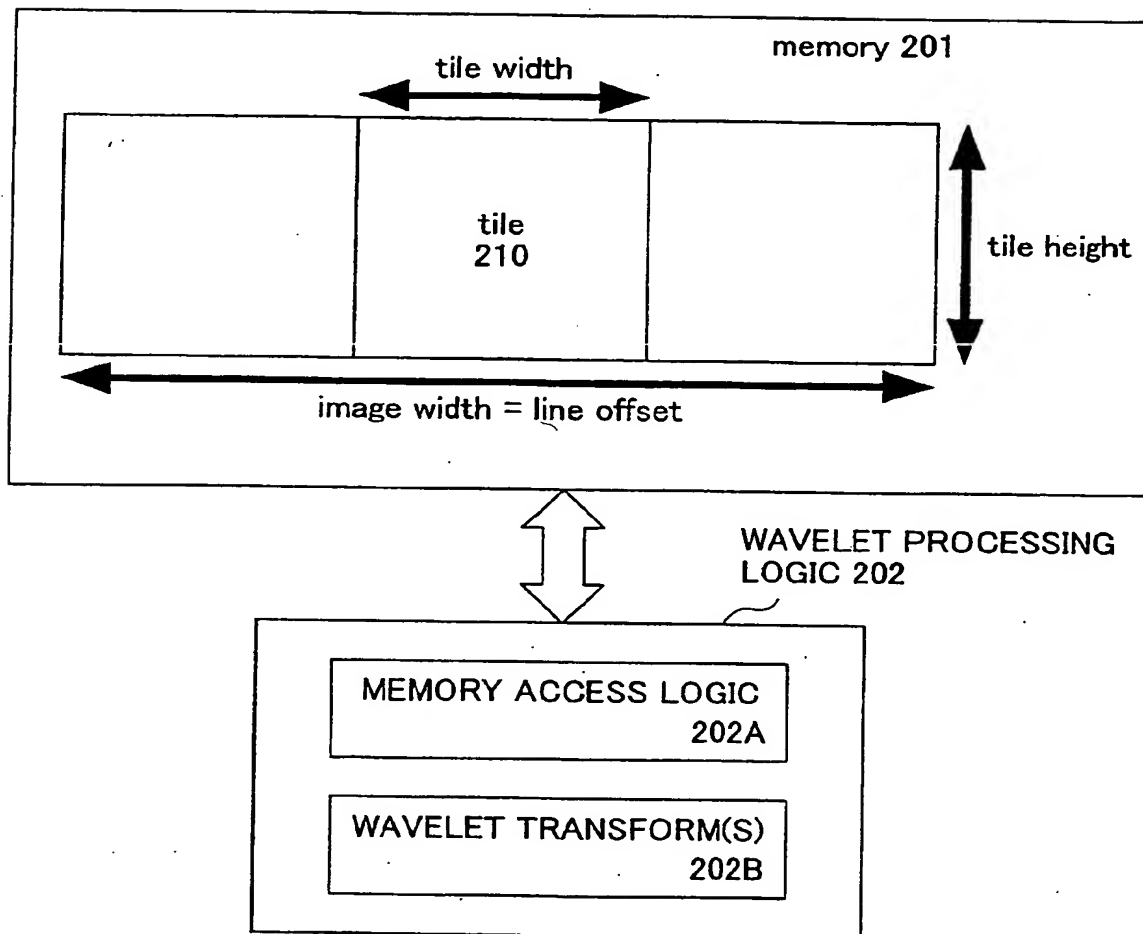



FIG.3A

FIG.3B

FIG.3C






FIG.4A




FIG.4B

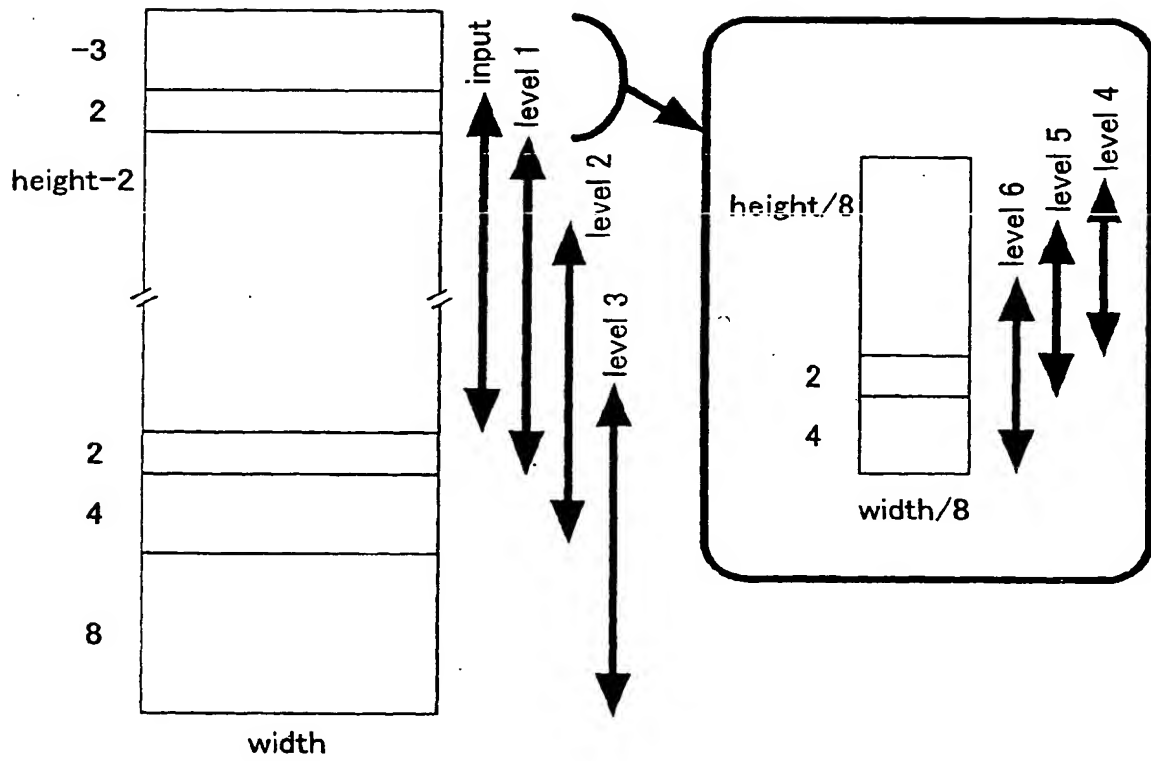


FIG.5

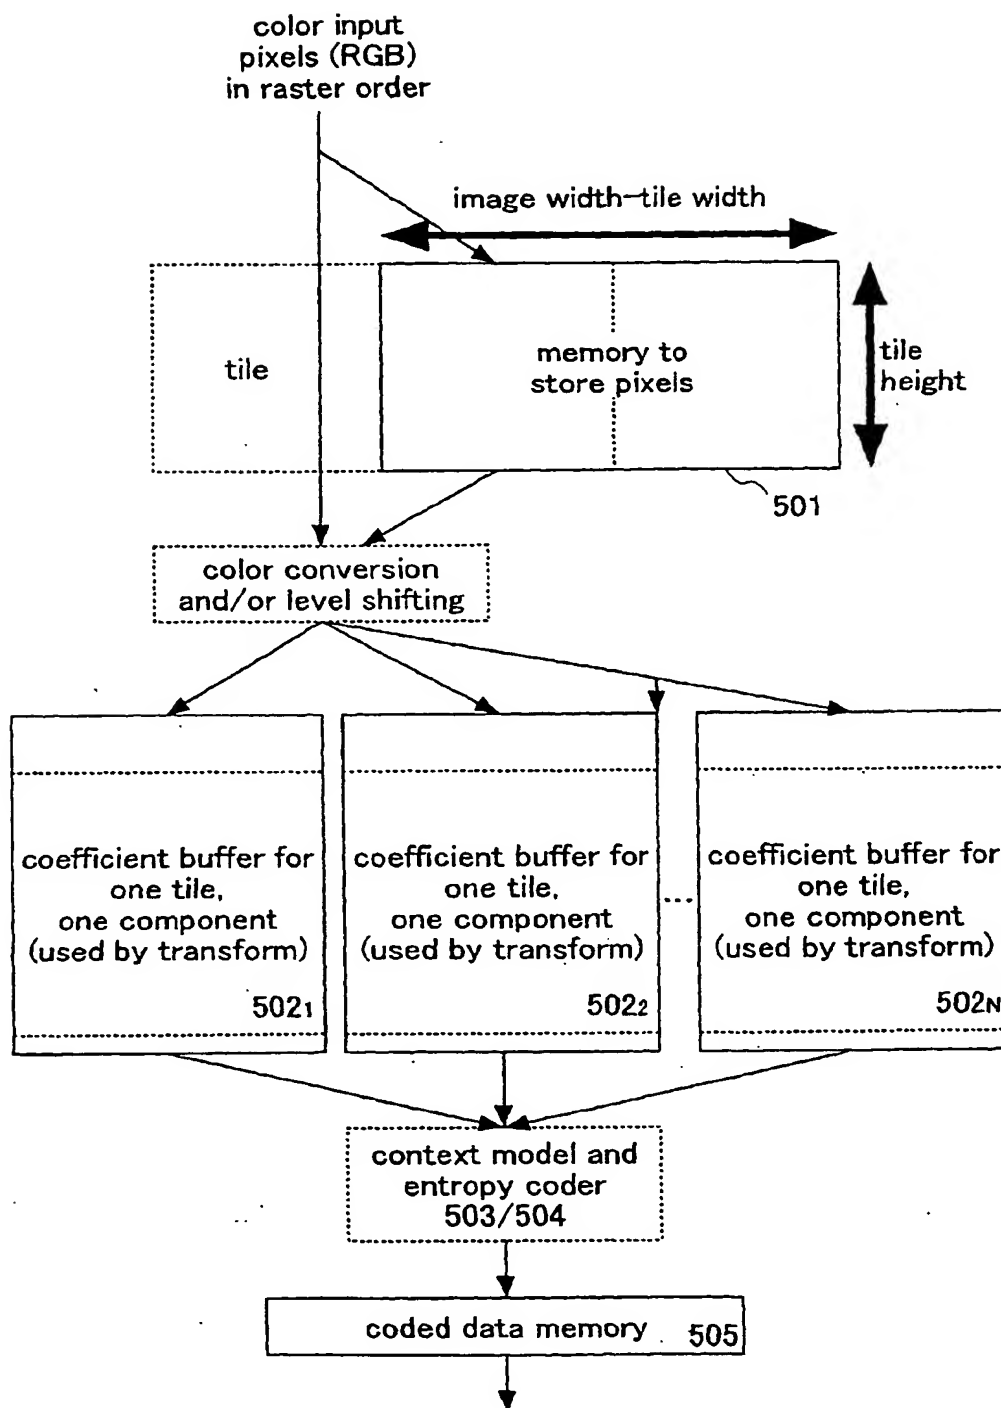


FIG.6A




FIG.6B

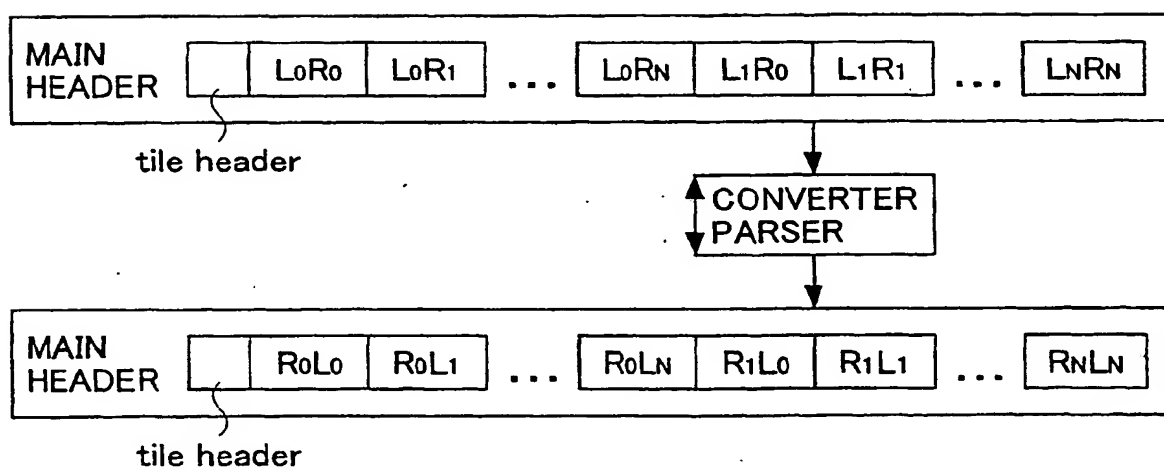


FIG.7A




FIG.7B




FIG.8

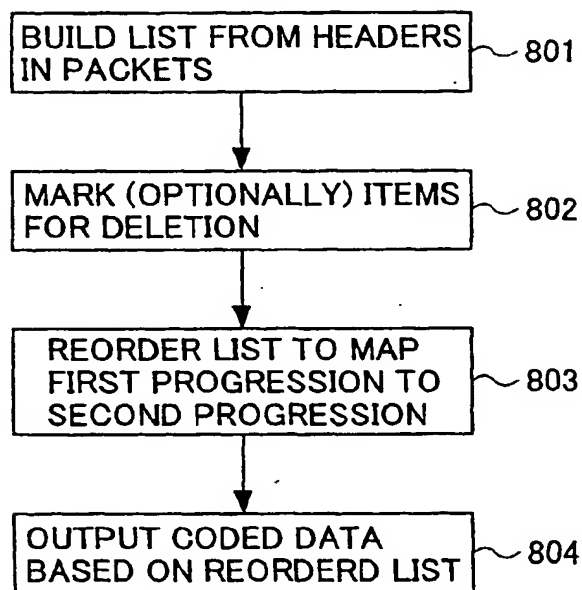


FIG.9

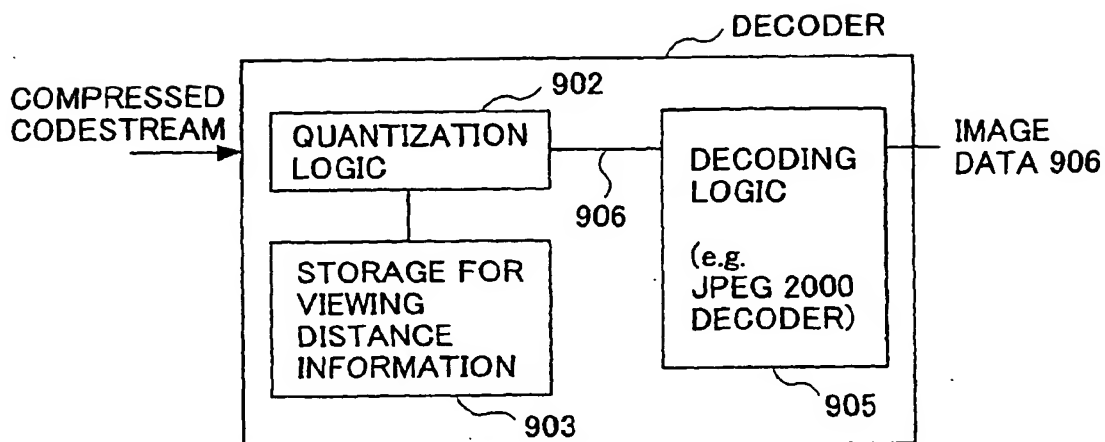


FIG.10




FIG.11

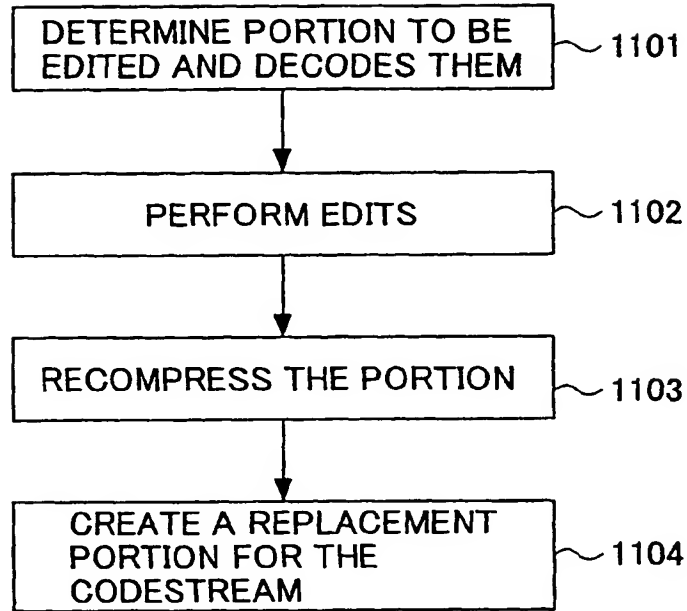


FIG.12




FIG.13




FIG.14




FIG.15A

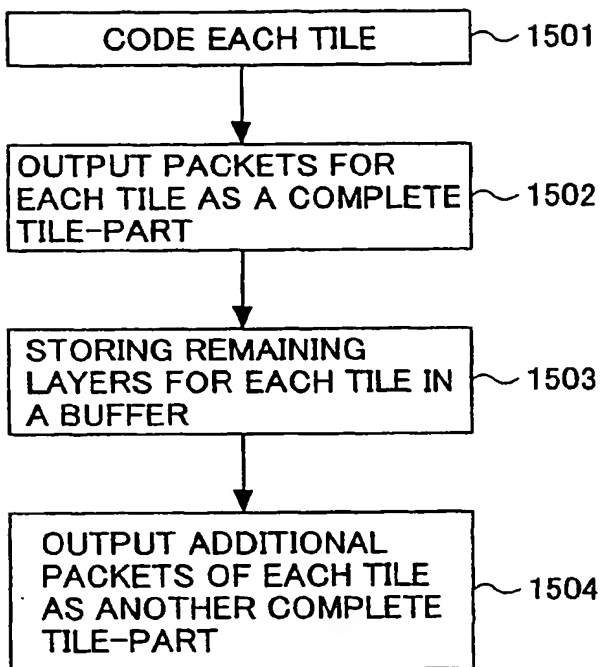


FIG.15B




FIG.16

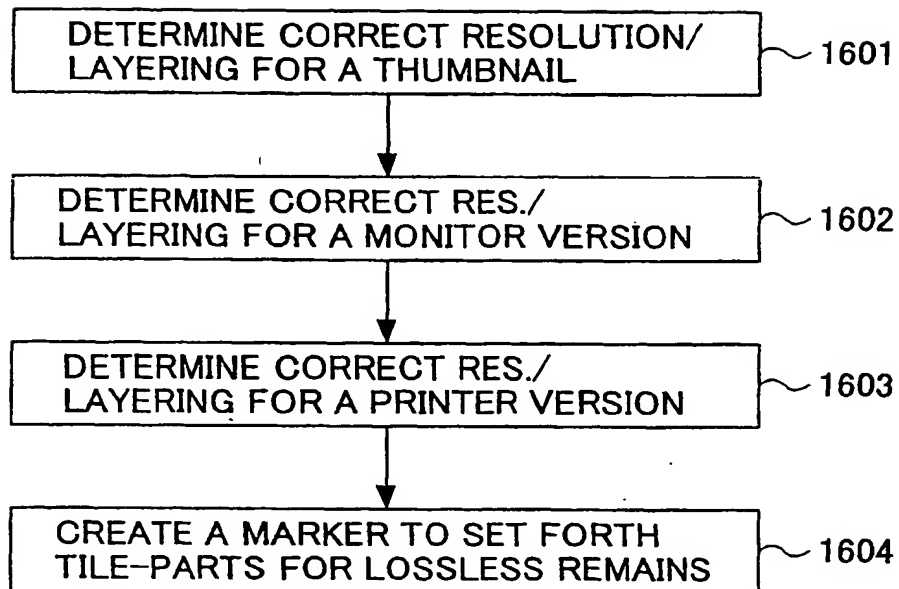


FIG.17

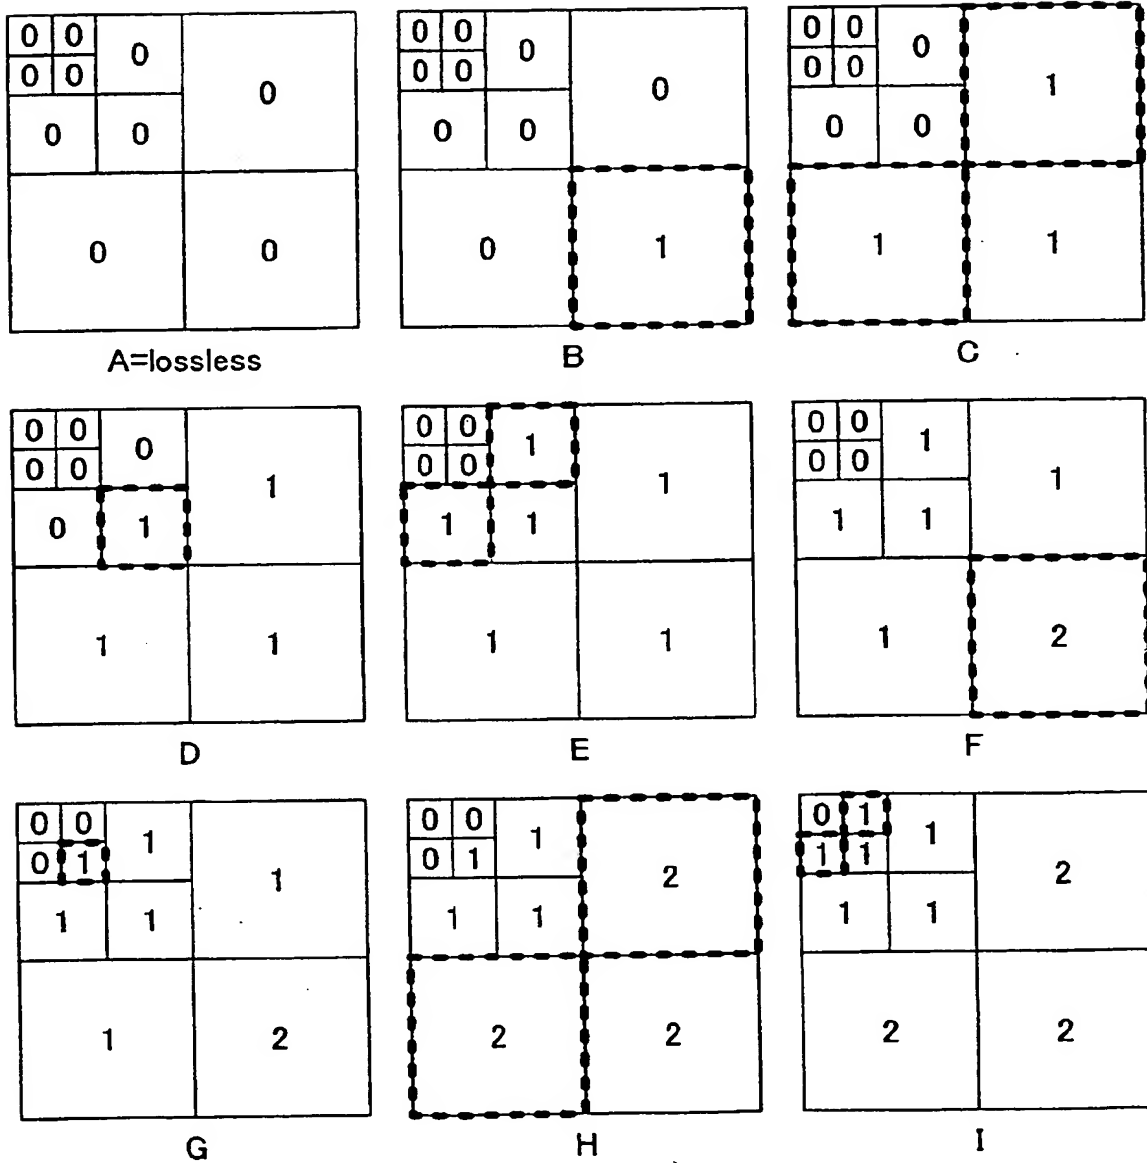


FIG.18




FIG.19




FIG.20

2000

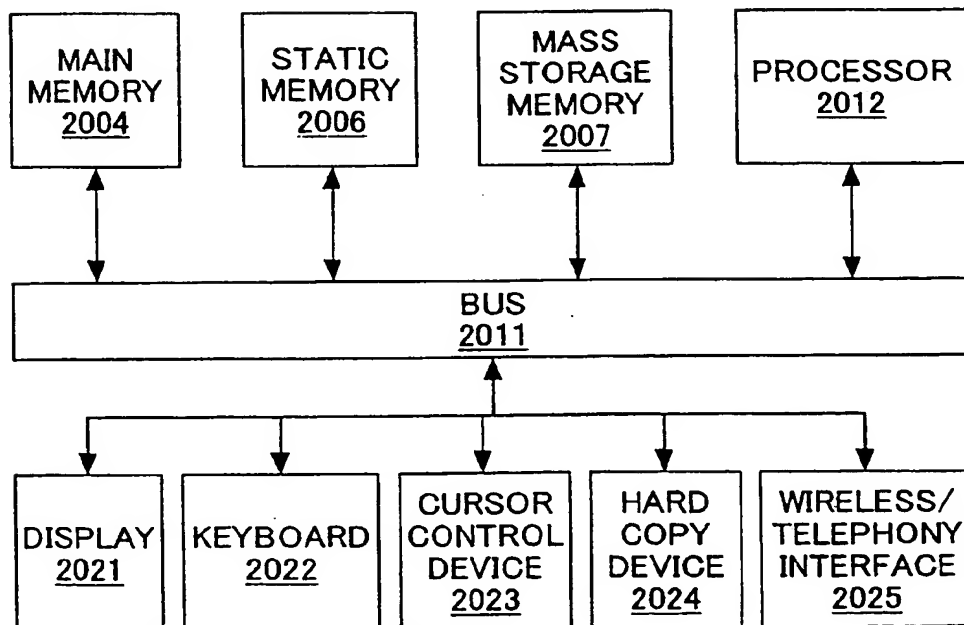


FIG.21

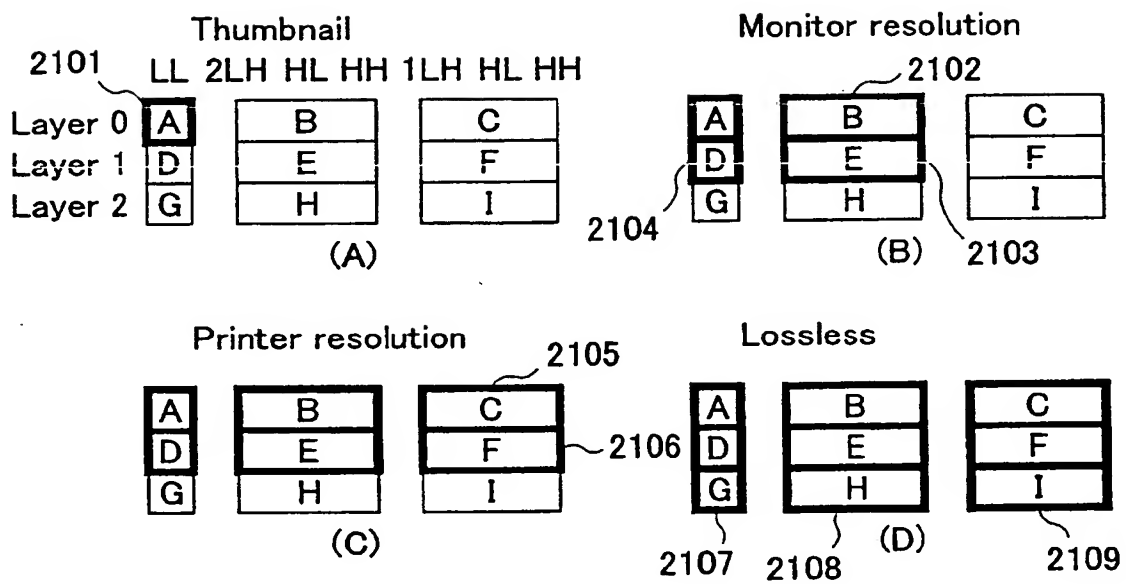


FIG.24




FIG.25

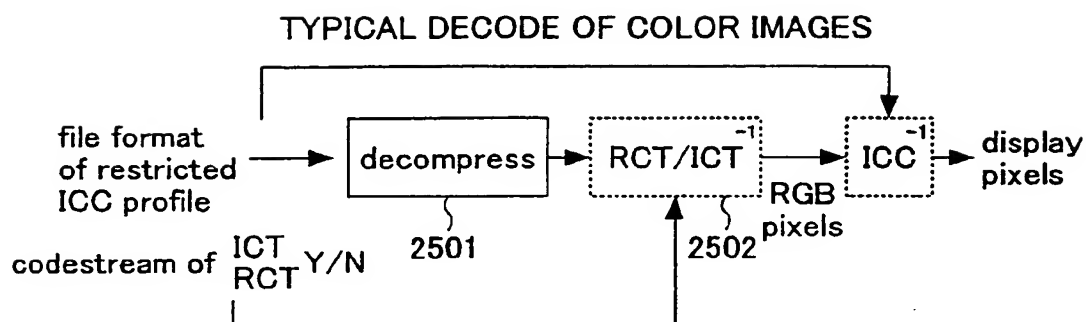


FIG.26




FIG.27





FIG.28





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 12 9242

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	CHRYSAFIS C ET AL: "Line based reduced memory, wavelet image compression" DATA COMPRESSION CONFERENCE, 1998. DCC '98. PROCEEDINGS SNOWBIRD, UT, USA 30 MARCH-1 APRIL 1998, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 30 March 1998 (1998-03-30), pages 398-407, XP010276575 ISBN: 0-8186-8406-2 * sections 1, 2 *	1-22	H04N7/26
X	GORDON B M ET AL: "A 1.2 MW VIDEO-RATE 2-D COLOR SUBBAND DECODER" IEEE JOURNAL OF SOLID-STATE CIRCUITS, IEEE INC. NEW YORK, US, vol. 30, no. 12, 1 December 1995 (1995-12-01), pages 1510-1516, XP000557257 ISSN: 0018-9200 * section 11.C. * * figure 7 *	1-22	
X	MENG T H: "A wireless portable video-on-demand system" VLSI DESIGN, 1998. PROCEEDINGS., 1998 ELEVENTH INTERNATIONAL CONFERENCE ON CHENNAI, INDIA 4-7 JAN. 1998, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 4 January 1998 (1998-01-04), pages 4-9. XP010263397 ISBN: 0-8186-8224-8 * section 2.3. *	1-22	H04N
A	US 5 414 780 A (CARNAHAN SHAWN V A) 9 May 1995 (1995-05-09) * column 3, line 40 - column 5, line 10 * * column 6, line 1 - column 6, line 4 * * claims 1-4 *	1-22	
-/-			
The present search report has been drawn up for all claims			
Place of search		Date of completion of the search	Examiner
THE HAGUE		13 May 2002	Fassnacht, C
CATEGORY OF CITED DOCUMENTS		1: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date I: document cited in the application L: document cited for other reasons A: member of the same patent family, corresponding document	
X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document			

EP FORM 1503 32.92 (04/2001)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 12 9242

DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim
A	EP 1 035 511 A (CANON KK) 13 September 2000 (2000-09-13) * paragraph [0003] - paragraph [0006] * * paragraph [0021] - paragraph [0038] * * figures 2-8 *	1-22
A	US 5 966 465 A (GORMISH MICHAEL J ET AL) 12 October 1999 (1999-10-12) * column 3, line 23 - column 3, line 30 * * column 28, line 54 - column 30, line 40 * * figures 16A,16B,18 *	1-22
A	GB 2 303 030 A (RICOH KK) 5 February 1997 (1997-02-05) * page 62, line 13 - page 67, line 9 *	1-22
		TECHNICAL FIELDS SEARCHED (Int.Cl.7)
The present search report has been drawn up for all claims		
Place of search	Date of completion of the search	Examiner
THE HAGUE	13 May 2002	Fassnacht, C
CATEGORY OF CITED DOCUMENTS		
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background Q : non-written disclosure P : intermediate document 1 : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application I : document cited for other reasons 8 : member of the same patent family, corresponding document		

EP FORM 1503 (2.8.2) (F04001)

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO.

EP 01 12 9242

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

13-05-2002

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
US 5414780	A	09-05-1995	JP	7184201 A	21-07-1995
EP 1035511	A	13-09-2000	AU	2080600 A	05-10-2000
			EP	1035511 A2	13-09-2000
			JP	2000308060 A	02-11-2000
US 5966465	A	12-10-1999	US	5867602 A	02-02-1999
			US	5748786 A	05-05-1998
			US	6229927 B1	08-05-2001
			US	6141446 A	31-10-2000
			DE	19626600 A1	09-01-1997
			GB	2303031 A ,B	05-02-1997
			GB	2313757 A ,B	03-12-1997
			JP	9027912 A	28-01-1997
			US	5881176 A	09-03-1999
			DE	19534943 A1	28-03-1996
			FR	2724791 A1	22-03-1996
			FR	2724792 A1	22-03-1996
			GB	2293733 A ,B	03-04-1996
			GB	2293734 A ,B	03-04-1996
			GB	2302244 A ,B	08-01-1997
			GB	2302245 A ,B	08-01-1997
			GB	2302488 A ,B	15-01-1997
			GB	2305829 A ,B	16-04-1997
			JP	8139935 A	31-05-1996
			NL	1001248 C2	04-06-1997
			NL	1001248 A1	21-03-1996
			US	6195465 B1	27-02-2001
			US	6222941 B1	24-04-2001
GB 2303030	A	05-02-1997	US	6195465 B1	27-02-2001
			CN	1145553 A	19-03-1997
			DE	19626615 A1	09-01-1997
			JP	9027752 A	28-01-1997

EPO FORM P0159

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82